EPRI | ELECTRIC POWER RESEARCH INSTITUTE

EGEAS Capabilities Manual

Version 13.1

3002014878

EGEAS Capabilities Manual

Version 13.1

3002014878

Software Manual, June 2020

EPRI Project Manager

A. Diamant

ELECTRIC POWER RESEARCH INSTITUTE
3420 Hillview Avenue, Palo Alto, California 94304-1338 ▪ PO Box 10412, Palo Alto, California 94303-0813 ▪ USA
800.313.3774 ▪ 650.855.2121 ▪ askepri@epri.com ▪ www.epri.com

DISCLAIMER OF WARRANTIES AND LIMITATION OF LIABILITIES

ELECTRIC POWER RESEARCH INSTITUTE, INC. (“EPRI”) RESERVES ALL RIGHTS IN THE PROGRAM AS DELIVERED. THE PROGRAM OR ANY PORTION THEREOF MAY NOT BE REPRODUCED IN ANY FORM WHATSOEVER EXCEPT AS PROVIDED BY LICENSE, WITHOUT THE CONSENT OF EPRI.

A LICENSE UNDER EPRI'S RIGHTS IN THE PROGRAM CAN BE OBTAINED DIRECTLY FROM EPRI.

THE EMBODIMENTS OF THIS PROGRAM AND SUPPORTING MATERIALS MAY BE INDEPENDENTLY AVAILABLE FROM ELECTRIC POWER SOFTWARE CENTER (EPSC) FOR AN APPROPRIATE DISTRIBUTION FEE.

ELECTRIC POWER SOFTWARE CENTER (EPSC)
9625 RESEARCH DRIVE
CHARLOTTE, NC 28262

THIS NOTICE MAY NOT BE REMOVED FROM THE PROGRAM BY ANY USER THEREOF.

NEITHER EPRI, ANY MEMBER OF EPRI, THE ORGANIZATION(S) BELOW, NOR ANY PERSON ACTING ON BEHALF OF ANY OF THEM:

1. MAKES ANY WARRANTY OR REPRESENTATION WHATSOEVER, EXPRESS OR IMPLIED, INCLUDING ANY WARRANTY OF MERCHANTABILITY OR FITNESS OF ANY PURPOSE WITH RESPECT TO THE PROGRAM ; OR
2. ASSUMES ANY LIABILITY WHATSOEVER WITH RESPECT TO ANY USE OF THE PROGRAM OR ANY PORTION THEREOF OR WITH RESPECT TO ANY DAMAGES WHICH MAY RESULT FROM SUCH USE.

RESTRICTED RIGHTS LEGEND: USE, DUPLICATION, OR DISCLOSURE BY THE GOVERNMENT IS SUBJECT TO RESTRICTION AS SET FORTH IN PARAGRAPH (G) (3) (I), WITH THE EXCEPTION OF PARAGRAPH (G) (3) (I) (B) (5), OF THE RIGHTS IN TECHNICAL DATA AND COMPUTER SOFTWARE CLAUSE IN FAR 52.227-14, ALTERNATE III.

REFERENCE HEREIN TO ANY SPECIFIC COMMERCIAL PRODUCT, PROCESS, OR SERVICE BY ITS TRADE NAME, TRADEMARK, MANUFACTURER, OR OTHERWISE, DOES NOT NECESSARILY CONSTITUTE OR IMPLY ITS ENDORSEMENT, RECOMMENDATION, OR FAVORING BY EPRI.

THE FOLLOWING ORGANIZATION, UNDER CONTRACT TO EPRI, PREPARED THIS REPORT:

NGPlanning LLC

NOTE

For further information about EPRI, call the EPRI Customer Assistance Center at 800.313.3774 or e-mail askepri@epri.com.

Electric Power Research Institute, EPRI, and TOGETHER…SHAPING THE FUTURE OF ELECTRICITY are registered service marks of the Electric Power Research Institute, Inc.

Copyright © 2020 Electric Power Research Institute, Inc. All rights reserved.

ACKNOWLEDGMENTS

The following organization, under contract to the Electric Power Research Institute (EPRI), prepared this report:

NGPlanning LLC
4 Cambridge Road
Holmdel, NJ 07733

Principal Investigators
W.S. Ng
W.D. Fleck

This report describes research sponsored by EPRI.

This publication is a corporate document that should be cited in the literature in the following manner:

EGEAS Capabilities Manual: Version 13.1. EPRI, Palo Alto, CA: 2020. 3002014878.

PRODUCT DESCRIPTION

Introduction

EGEAS, an EPRI product, developed in the early 1980s was the fore-runner of the current generation planning simulation and optimization models in the market and is still being used by a number of utilities and regional planning organizations. It currently runs on PC platform with the following operating systems: Microsoft Windows 7, and Windows 8 (32-bit and 64-bit).

Background

EGEAS is a state-of-the-art modular production costing and generation expansion software package developed under EPRI sponsorship for use by utility planners develop and to evaluate integrated resource plans, avoided costs, and develop plant life management plans. It also has modules which accommodate demand-side management options and facilitate development of environmental compliance plans to meet mandated state and federal regulations. EGEAS contains three capacity optimization options that range from preliminary analysis tools based on screening curves to sophisticated non-linear optimization using generalized Benders decomposition and dynamic programming algorithms. Detailed probabilistic production costing algorithm is also available for production cost calculations and reliability analysis.

Key Features

Dynamic programming algorithm to develop candidate portfolios from identified alternatives meeting a reliability constraint (e.g., maintain 10% capacity reserve margin).

Conducts minimum present value revenue requirements (PVRR) or lowest electric rate economic ranking of candidate portfolios dispatched with an existing and future set of assets.

Once data inputs are set-up, many different scenarios can be run to test different generation plans very easily and quickly.

Simple economics based on present worth method - easy to understand and explain results.

Used by many utilities and various regional planning organizations, and results are understood and well accepted by regulators.

Current Applications

Integrated Resource Planning (IRP) studies.

Regional Resource Forecasting (RRF) studies.

Front-end screening tool for future resources.

Existing asset retirement evaluations.

Emission evaluations from power plants.

Development of compliance plans to meet environmental regulations.

What's new?

The software originally was developed in FORTRAN to run on mainframe computers in a batch mode, and later was converted to run on laptop computers. In the 1990s, Stone &Webster Management Consultants Inc. (SWMCI), a software licensee, developed a Graphical User Interface (GUI) to make the data input easier. In 2013, EPRI acquired legal rights to the EGEAS

GUI. Version 10.0 of the EGEAS software package released in 2014 incorporated version 9.02 with the integrated EGEAS GUI.

In 2015, EPRI contracted with NG Planning LLC to develop EGEAS version 11.0 which incorporates two new enhancements: (i) an option for users to apply a penalty factor to the "dump energy" cost; and, (ii) the ability for users to specify Renewable Portfolio Standard (RPS) constraints as part of the optimization of the least-cost generation expansion plan.

EGEAS Version 12 adds constraints on emission rates, provides generating unit capacity reports, and greatly increases limits on many quantities. EGEAS Version 13 adds improvements in modeling of renewable resources, provision for input of load and generation data at sub-hourly time intervals, and several other minor enhancements. EGEAS Version 13.1 corrects a few errors, but adds no new functionality.

Keywords

EGEAS
IRP
Resource optimization
Supply-side & demand-side resources
Production costs
System reliability
Environmental compliance
Avoided costs
Marginal costs
Incremental costs
Fuel costs
Capacity costs.

ABSTRACT

EGEAS, an EPRI product, developed in the early 1980s was the fore-runner of the current generation planning simulation and optimization models in the market and is still being used by a number of utilities and regional planning organizations. It currently runs on PC platform with the following operating systems: Microsoft Windows 7, and Windows 8 (32-bit and 64-bit).

EGEAS is a state-of-the-art modular production costing and generation expansion software package developed under EPRI sponsorship for use by utility planners develop and to evaluate integrated resource plans, avoided costs, and develop plant life management plans. It also has modules which accommodate demand-side management options and facilitate development of environmental compliance plans to meet mandated state and federal regulations. EGEAS contains three capacity optimization options that range from preliminary analysis tools based on screening curves to sophisticated non-linear optimization using generalized Benders decomposition and dynamic programming algorithms. Detailed probabilistic production costing algorithm is also available for production cost calculations and reliability analysis.

PREFACE

The Electric Generation Expansion Analysis System (EGEAS) is a modular state-of-the-art generation expansion software package. EGEAS is used by utility planners to produce integrated resource plans, evaluate independent power producers, develop avoided costs and environmental compliance plans, and analyze life extension alternatives. EGEAS is a set of computer modules that determine an optimum expansion plan or simulate detailed production costs for a prespecified plan. Expansion plans are defined by the type, size, and installation date for each new generating facility or demand-side management resource. Optimum expansion plans are developed in terms of annual costs, operating expenses, and carrying charges on investment; average system costs and financial ratios. The objective is to find an integrated resource plan that meets the objective function specified by you. Typical objective functions include: minimizing total costs, minimizing customer rates, minimizing societal costs, maximizing earnings.

EGEAS is designed to allow you to determine the amount of detail to include in each study. Within the expansion algorithms, EGEAS contains several expansion plan analysis options, ranging from preliminary analysis tools based on screening curves to sophisticated non-linear analysis tools using a Generalized Benders Decomposition algorithm or a Dynamic Programming algorithm. A stand-alone, detailed probabilistic production-costing algorithm is also available for production costing and reliability analysis.

EGEAS can handle a wide variety of generation technologies including thermal (nuclear, fossil, combined cycle, combustion turbine), limited energy (hydroelectric, interruptible rates), storage (pumped hydro, cool storage batteries, compressed air), and non-dispatchable technologies (solar, wind, cogeneration, conservation load management). Additional features include interconnections to neighboring utilities or power pools, purchase and sale contracts, environmental constraints and calculations and automatic sensitivity analysis. EGEAS can also analyze DSM options such as conservation, strategic marketing, load management, storage, and rate design.

The EGEAS Capabilities Manual describes, in detail, specific features of the EGEAS program. This capabilities manual includes the following chapters:

Introduction

Expansion Plan Simulation

Production Costing

Optimization Analysis

Benders' Decomposition

Dynamic Programming

Financial Analysis

System Interconnections

Environmental Compliance Planning

Additional Capabilities

Bibliography

For specific instructions on using EGEAS refer to the EGEAS User's Guide: Version 13.1.

CONTENTS

1 INTRODUCTION 1-1
Overview of EGEAS 1-1
This Manual 1-2
Abbreviations 1-3
Additional Reading 1-4
2 EXPANSION PLAN SIMULATION 2-1
Overview 2-1
Economic Assumptions 2-1
Periods of Analysis 2-2
Periods within a Year 2-2
System Load 2-3
Load Duration Curves 2-3
Customer Classes 2-6
Interconnected Systems 2-6
Input Requirements 2-7
Methodology 2-7
Resources: Supply-Side and Demand-Side 2-12
Resource Type and Status 2-12
Capacity 2-13
Heat Rate 2-14
Forced Outage Rate 2-17
Maintenance 2-19
Fuel 2-19
Limited Energy and Storage Plants 2-20
Non-Dispatchable Technology Resources 2-20
Expansion Plan Costs 2-21
Capital Fixed Charges 2-21
Operating Costs 2-22
Detailed Costs 2-23
Purchase Costs and Sale Revenues 2-23
Unserved Energy Cost 2-24
Emission Costs 2-24
Reliability Measures 2-24
Reserve Margin 2-24
Loss of Load Probability 2-25
Relative Reliability 2-26
Unserved Energy 2-26
End Effects 2-26
Input Requirements 2-28

System Load 2-28
Existing Resources 2-28
New Resources 2-29
Modeling Parameters 2-29
3 PRODUCTION COSTING 3-1
Overview 3-1
Load Model 3-1
Resources Model 3-2
Probabilistic Production Costing Algorithm 3-3
Single-Block Thermal Resources 3-3
Multiple-Block Thermal Resources 3-7
Dispatch Order 3-9
Loading Blocks 3-10
Must-Run Units 3-10
Spinning Reserve 3-12
Dispatch Modifiers 3-15
Limited Energy Plants 3-16
Model Description 3-16
Methodology 3-16
Storage Units 3-21
Model Description 3-21
Short-Term Decision 3-23
Medium-Term Decision 3-25
Non-Dispatchable Technologies 3-29
Model Description 3-29
Methodology 3-30
Purchase and Sale Contracts 3-31
Input Requirements 3-32
Methodology 3-33
Maintenance Scheduling Options 3-34
Maintenance Cycles 3-34
Maintenance Scheduling Algorithm 3-35
Dual Fuels and Fuel Constraints 3-36
Input Requirements 3-37
Methodology 3-38
Representation of Load Duration Curve 3-41
Piecewise Linear Representation 3-41
Analytical Representation With Cumulants 3-42
Comparison of Representations 3-44
4 OPTIMIZATION ANALYSIS 4-1
Overview 4-1

Objective Functions....4-2
Screening Tests for DSM....4-3
Total Resource Cost Test (TRC)....4-3
Rate Impact Measure....4-3
No-Losers Test....4-4
Societal Test....4-4
Decision Variables....4-4
Constraints....4-5
Reliability Constraints....4-5
Tunnel Constraints....4-5
Financial Constraints....4-5
Fuel Constraints....4-6
Emission Constraints....4-6
Other Constraints....4-6
Optimization Methods....4-6
Screening Curves....4-9
Pre-Specified Pathway....4-12
5 BENDERS' DECOMPOSITION....5-1
Overview....5-1
Objective Function....5-1
Decision Variables....5-1
Constraints....5-2
Methodology....5-2
Cost Constraints....5-3
Unserved Energy Constraints....5-7
Subproblem....5-8
Master Problem....5-9
Convergence....5-10
Limitations....5-11
6 DYNAMIC PROGRAMMING....6-1
Overview....6-1
Objective Function....6-1
Decision Variables....6-1
Constraints....6-2
Methodology....6-3
State Generation....6-5
State Evaluation....6-7
State Compression....6-8
Forward Dynamic Programming....6-9
Backward Dynamic Programming....6-13
Construction of Expansion Plans....6-16

Simulation of Expansion Plans 6-17
Installation Year Estimation 6-17
Suboptimal Plans 6-18
Dependency Constraints 6-19
Prerequisite Alternatives 6-19
Retirement of an Existing Resource 6-22
Replacement of an Existing Unit 6-23
Mutually-Exclusive Alternatives 6-23
Single-Year Alternatives 6-23
Multiple-Unit Capital Costs 6-23
Cost Postprocessor 6-24
Limitations 6-24
Practical Guidelines 6-25
7 FINANCIAL ANALYSIS 7-1
Overview 7-1
Format 1 – Levelized Fixed Charges 7-1
Calculations 7-2
Format 2 – Annual Fixed Charges 7-2
Capital Cost Structure by Resource 7-3
Rate of Return on Rate Base 7-3
Fixed Charges Calculations 7-4
Format 3 – Construction Costs 7-6
Separating Equity and Debt Portions of AFUDC 7-8
Construction Cost Calculations 7-8
Capital Expenditures Extended Beyond Installation Year 7-13
Sample Calculation 7-13
Nature of Financial Constraints 7-19
Profitability 7-19
Liquidity 7-19
Price Constraints 7-20
Feasibility Measures 7-20
Non-EGEAS Assets 7-21
Total Rate Base and CWIP 7-22
Earning Assets Ratio 7-24
Interest Coverage Ratio 7-25
System Average Cost 7-26
Customer Revenue Accounting 7-28
8 SYSTEM INTERCONNECTIONS 8-1
Overview 8-1
Reserve Sharing 8-1
Method 1 – Tie-Line Capacity Limit 8-1

Method 2 – Percent Load Limit 8-2
Method 3 – Multi-State Unit 8-2
Economy Interchange 8-2
Methodology 8-3
Tie-Line Limits 8-4
Cost Accounting 8-5
Input Requirements 8-10
9 ENVIRONMENTAL COMPLIANCE PLANNING 9-1
Overview 9-1
Plant Emissions 9-2
Generic Sites 9-3
Emission Limits 9-3
Allowance Costs and Credits 9-4
Unit Emission Constraints 9-5
Site Emission Constraints 9-5
System Emission Dispatch and Constraint Options 9-6
Banking of Emission Allowances 9-7
Fuel Switching 9-8
Emissions Effects 9-9
Emission Reduction Constraints 9-9
Overview 9-9
Emission Rate Calculation 9-9
Constraint Methodology 9-11
Adjust Dispatch Method 9-11
Reject State Method 9-12
Annual vs. Segment-Level Constraints 9-12
Multiple Constraints 9-13
Practical Considerations and Warnings 9-14
Limitations 9-15
10 ADDITIONAL CAPABILITIES 10-1
Overview 10-1
Uncertainties Modeling 10-1
Automatic Sensitivity Analysis 10-1
Avoided Costs 10-2
Avoided Capacity Costs 10-3
Avoided Energy Costs 10-5
Hourly Avoided Capacity Cost File 10-5
Incremental Costs 10-6
Hourly Incremental Cost Files 10-8
Price Elasticity and Feedback Mechanisms 10-9
Bid-Based Pricing 10-11

Limitations 10-14
Renewable Portfolio Standard (RPS) Constraints 10-14
RPS Calculation 10-14
RPS Constraint Application 10-15

LIST OF FIGURES

Figure 2-1 Subperiod Definition Example 2-4
Figure 2-2 Conversion of Time-Dependent Curve to Load Duration Curve 2-5
Figure 2-3 Deterministic Operating Scheduling 2-6
Figure 2-4 System and Class Load Curves 2-7
Figure 2-5 Heat Input-Output Curve 2-16
Figure 3-1 Conversion of Load Duration Curve to Cumulative Probability Curve 3-2
Figure 3-2 Loading Resources Without Outages 3-4
Figure 3-3 Graphic Illustration of Convolution 3-6
Figure 3-4 LOLP and Unserved Energy 3-7
Figure 3-5 Trial Loading of an LEP 3-18
Figure 3-6 Loading of a Delayed LEP 3-19
Figure 3-7 Delayed LEP Loaded at Reduced Capacity 3-20
Figure 3-8 Dispatch of a Storage Charging Side 3-24
Figure 4-1 Screening Curves 4-13
Figure 5-1 Benders Master Problem and Subproblem 5-4
Figure 5-2 Expansion Plant Cost Components 5-5
Figure 5-3 Benders Cost Constraints 5-6
Figure 5-4 Benders Unserved Energy Constraint 5-8
Figure 6-1 Dynamic Program States 6-4
Figure 6-2 Committed Costs and Backward Pointers 6-12
Figure 6-3 Costs to Go and Forward Pointers 6-15

LIST OF TABLES

Table 4-1 Comparison of Optimization Methods....4-7
Table 6-1 Expansion Plans from Forward DP 6-13
Table 6-2 Expansion Plans from Forward and Backward DP 6-16
Table 8-1 Input Requirements for Interconnections....8-10
Table 10-1 Incremental Cost Table 10-8

1
INTRODUCTION

Overview of EGEAS

The primary purpose of EGEAS is to find the best possible integrated resource plan for meeting increasing demand by expanding an electric system's generating capacity or reducing load via demand-side management. Optimized plans may be derived for each sensitivity or scenario case and for each objective function specified. In addition, EGEAS includes enhanced production costing capabilities that can produce a detailed simulation of the specified expansion plans.

EGEAS includes two rigorous optimization tools: Generalized Benders' Decomposition is a non-linear technique based on an iterative interaction between a linear master problem and a non-linear probabilistic production costing subproblem. The dynamic program (DP) is based on an enumeration of possible capacity additions and a standard dynamic programming technique for selecting the optimal path (set of installation steps).

Each of the optimization tools supports constraints that limit the range of plans. Reliability constraints eliminate plans with insufficient or over abundant reserve capacity, energy, or loss-of-load probability. Economic constraints eliminate plans that are economically infeasible (low earning assets ratio, low interest coverage ratio, or a large increase in system average rate). Tunnel constraints specify upper and lower limits on the cumulative number of resource or year-by-year additions, of a particular resource type, installed up through a given year; they are used to guide the optimization process away from plans that are clearly not optimal or not feasible, and thus reduce the EGEAS run time. The DP program also supports environmental and fuel availability constraints.

EGEAS has two other tools for analyzing capacity planning alternatives. The screening curve option can be used to sift through large numbers of resources and eliminate those that are clearly not optimal. The pre-specified pathway option provides more detailed analysis of a plan than is computationally feasible within an optimization.

The expansion plans analyzed by EGEAS can include a wide variety of technologies for electricity generation and demand-side management (DSM). Among the supply-side options are thermal units (i.e., nuclear, fossil steam, combined cycle, combustion turbine), retirement of existing facilities, staged resources (i.e., IGCC), life extension programs, independent power producers, hydroelectric units (or limited energy), storage (i.e., pumped hydro, batteries, compressed air, superconducting magnets), and non-dispatchable (hourly) technologies (i.e., photovoltaics, wind, and cogeneration). Among the DSM options are conservation, strategic marketing, load management, storage and rate designs, each of which may have several hundred DSM technologies that can be considered as resource options.

Additional modeling capabilities of EGEAS include purchase and sale contracts, interconnections with other systems, incremental (avoided) capacity and operating costs, customer class revenues and sales, and environmental calculations. EGEAS can also perform automatic sensitivity analyses to test the sensitivity of results to variations in the input assumptions.

The production costing capabilities in EGEAS include:

four levels of capacity (rated, operating, emergency, and reserves),
changing capacity levels by year and month,
loading points (up to 5) for capacities, heat rates, and forced outages,
automatic maintenance scheduling,
spinning reserve designations and options,
monthly fuel limitations,
fuel targets (minimum, maximum and target percentages),
costs (fixed O&M, variable O&M, and T&D costs),
dispatch modifier costs, and
limited energy data (by year and by month),

For DSM resources additional capabilities include:

customer costs,
rebound benefits,
direct customer benefits,
benefits related to rate changes,
transmission and distribution costs and savings,
price elasticities by customer class, and
customer class rates (demand, energy and customer charges).

This Manual

This Capabilities Manual describes the types of analyses EGEAS performs and the methods it uses. For each major area of analysis, this manual describes the basic capabilities and major options, outlines the methodology, presents the assumptions and limitations, and discusses the interactions with other analysis features. The following chapters describe the specific applications of EGEAS in more detail:

Chapter 2 provides a general background of all the important information including definitions and assumptions that EGEAS uses in simulating expansion plans.

Chapter 3 presents the basic structure and several extended modeling capabilities of the production costing algorithms used in EGEAS.

Chapter 4 discusses the different techniques EGEAS uses for analyzing expansion plans. While this chapter describes the primary modeling capabilities in all of the four analysis tools available in EGEAS, the next two chapters discuss the two major optimization tools, namely Benders' decomposition and dynamic program, in greater depth.

Chapter 5 describes how EGEAS performs the expansion plan optimization based on the generalized Benders' decomposition method.

Chapter 6 presents the dynamic programming methodology and all of the extended modeling capabilities EGEAS provides with the dynamic program.

Chapter 7 documents the financial calculations that EGEAS performs, in details, including equations and sample calculations.

Chapter 8 describes how EGEAS handles system interconnections in terms of reserve sharing and economy interchange.

Chapter 9 discusses the EGEAS capabilities in modeling plant emissions and emission constraints which allow you to evaluate environmental effects and various compliance strategies.

Chapter 10 describes several additional capabilities including uncertainty modeling, incremental and avoided cost calculations, and price elasticity and feedback mechanisms.

Abbreviations

The following abbreviations are used in this manual.

AFUDC	allowance for funds used during construction
Btu	British thermal unit
CWIP	construction work in progress
DP	dynamic program
DSM	demand-side management
EBIT	earnings before interest and taxes
GW	gigawatt = billion (10^9) watts
GWh	gigawatt hour = billion (10^9) watt hours
HR	hour
kW	kilowatt = thousand (10^3) watts
kWh	kilowatt hour = thousand (10^3) watt hours
K$	thousands of dollars
LDC	load duration curve
LEP	limited energy plant
LOLH	loss of load hours
LOLP	loss of load probability
MBtu	million (10^6) Btu
MW	megawatt = million (10^6) watts
MWh	megawatt hour = million (10^6) watt hours
M$	millions of dollars
NDT	non-dispatchable technology
O&M	operation and maintenance
T&D	transmission and distribution

Additional Reading

Additional sources and references may assist you in learning more about EGEAS. Information on the history of the model and the theories which underlie its development can be found in the following:

> *Electric Generation Expansion Analysis System, Volume 1: Solution Techniques, Computing Methods, and Results*, EPRI EL-2561, Volume 1, Project 1529-1, Final Report, August, 1982
>
> *Electric Generation Expansion Analysis System, Volume 2: Details of Solution Techniques, Data of Tests Systems, and Glossary of Terms*, EPRI EL-2561, Volume 2, Project 1529-1, Final Report, August, 1982 and in the Bibliography in Appendix A.

Information on operating instructions, input requirements and output results can be found in the EGEAS User's Guide.

2
EXPANSION PLAN SIMULATION

Overview

As its name implies, the Electric Generation Expansion Analysis System (EGEAS) analyzes expansion plans for an electric company. It can be used by investor-owned utilities (IOUs), independent power producers (IPPs), cooperatives, public power agencies and regional independent system operators.

An expansion plan is defined by a set of resources used to meet system demand. The specific expansion plan(s) to be simulated can be pre-specified by the EGEAS user, or determined by EGEAS using one of several optimization techniques. These optimization techniques are described in Chapters 4 through 6.

The remainder of this chapter includes definitions and assumptions that EGEAS uses in simulating expansion plans. Included are sections (i) how the system load is represented, and (ii) on the many types of supply-side and demand-side resources that can be used to meet the system load. Other sections describe supplementary features such as purchase and sale contracts, interconnections, and environmental issues. Later sections deal with expansion plan costs, reliability measures, and the treatment of end effects. The final section describes briefly the types of input data that you must supply.

Economic Assumptions

EGEAS uses a base year to serve as a reference point for economic data. The base year is usually the current year or some recent year. Most input cost data (e.g., installation costs and fuel costs) are assumed to be in base year dollars. EGEAS determines costs for later years by applying escalation rates to the base year values. Escalation rates may vary from year to year and different quantities can be escalated at different rates. In analyzing expansion plans, EGEAS determines costs for each year in current dollars (dollars for that year), then uses a discount rate to determine the present worth (value in base year dollars) of those costs. The discount rate measures the time value of money and is assumed to have the same value in every year. For customer costs used with demand-side management programs, EGEAS uses a separate customer discount rate. Inflation is assumed to be included in the discount rate and all escalation rates. A separate inflation rate can be specified in the model, but it is used only for the price elasticity calculation.

To illustrate the relationships among these quantities, consider the following example. Let the base year be 2010 and assume a discount rate of 15 percent. Assume further that a certain type of fuel has a cost in 2010 of $5 per MBTU, and that its cost can be expected to escalate at 10 percent per year in 2010 and 2011, and at 7 percent per year starting in 2012. The cost of one MBTU of fuel in 2015, in current dollars, is then:

$$(5.00) * (1.10)^2 * (1.07)^3 = 7.41$$

The present worth of this fuel is:

$$\frac{7.41}{(1.15)^5} = 3.68$$

Periods of Analysis

EGEAS performs generation expansion analysis over a period of years called the study period. The study period must start in or after the base year and may be one to 30 years in length. During this period, new resources are installed to meet forecasted load growth and to replace retiring units. EGEAS uses an extension period to model the end effects resulting from unused capital at the end of the study period. The extension period begins with the first year after the study period and may be finite or infinite in length. During the extension period, load remains constant at the same level as in the last study period year and no new units are installed. Any unit that retires during the extension period is assumed to be replaced with an identical unit. Costs continue to escalate at the prescribed rates. Extension periods are described more fully later in the End Effects section of this chapter.

Periods within a Year

The capabilities of resources can vary considerably over the course of a year. For example, the capacity of a thermal unit, or the amount of water for hydro generation, can vary seasonally. In addition, some resources (e.g., storage) are dispatched differently for different hours of the day or days of the week. However, the probabilistic production costing algorithm assumes that resources characteristics are constant. In cases where the variations are small, this limitation is not significant. In other cases, accuracy can be improved by dividing the year into smaller periods and performing production costing for each period with the proper values for unit characteristics. EGEAS can use four kinds of subperiods.

Each segment is a contiguous period consisting of an integral number of weeks; up to 13 segments per year are permitted. Examples of segments include seasons (13 weeks each) and months (four or five weeks each). Segments do not have to be equal in length. Generating capacities, energy limits, and fuel costs can be varied by segment. EGEAS uses segments to determine maintenance schedules. Note that EGEAS does not allow exact monthly periods of 30 or 31 days. They must be approximated with 4 and 5 week segments. However, EGEAS does provide a reporting option to scale the segment values to the lengths of individual months (e.g., multiply by 30/28 or 31/35).

Each subweek represents a portion of a week; up to 3 subweeks are allowed. Examples of subweek breakdowns include day versus night, weekday versus weekend, and weekday versus weeknight versus weekend. The number of subweeks must be the same for all segments, but the specific hours included in each subweek can be changed from segment to segment.

For improving the accuracy in modeling non-dispatchable technologies and correlation among loads from different systems for multi-area analysis, EGEAS provides the most detailed level of subperiod: time frames. Each time frame is a subset of the hours in a subweek, representing a smaller portion of a week. The maximum number of time frames, when summed over all subweeks, is 16. Note that time frames are used only for determining the impact of non-dispatchable technologies on the load and for determining the load for a pool dispatch in multi-

area modeling. The resulting loads are aggregated to the subweek level of detail for further processing

Each subyear is a group of one or more segments; up to four subyears are allowed. Subyears are useful when load and resource characteristics are already determined for individual segments, but a less detailed breakdown is desired to reduce computer run time. With subyears, EGEAS aggregates segment quantities to the subyear level before performing the analysis. Note that subyears can represent non-contiguous periods (e.g., by combining the first and last segments in a year); however, when you pre-specify maintenance periods, subyears should represent only contiguous groups of segments and they should be in calendar year order.

EGEAS uses subperiods as follows:

Hourly analysis: Modeling of purchase and sale contracts.

Time frame analysis: modeling of non-dispatchable technology units and determining the pool load for multi-area analysis.

Subweek analysis: Dispatch of thermal, hydro, and storage units for production costing; particularly important for storage operation.

Segment and subyear analysis: Dispatch of thermal, hydro, and storage units; allows seasonal variation in unit characteristics and maintenance scheduling.

Figure 2-1 illustrates an example of subperiods. The top section (a) shows the year divided into 12 segments of four and five weeks each, representing months. The segments are aggregated into three subyears, with subyear 2 representing summer (April through September) and subyears 1 (January through March) and 3 (October through December) representing winter. Subyears 1 and 3 are kept separate to preserve calendar year order for pre-specified maintenance scheduling.

The bottom section (b) of Figure 2-1 shows a one-week segment divided into three subweeks: weekday (1), weeknight (2), and weekend (3). Time frames (not shown) are similar to subweeks – each time frame includes a portion of one of the subweeks.

For budgeting purposes, you may want to have monthly reports. Under these circumstances you can request monthly data via the Report program. If this option is selected the program scales the segment information by the number of days in each month.

System Load

System demand is the sum of all customer loads (plus transmission and distribution losses) that must be met by the utility generating system being modeled. EGEAS seeks to meet the load using a combination of supply-side generating units, demand-side management resources, and purchases contracts. EGEAS does not include a load forecasting model, so you must specify the load for each year to be analyzed. However, the sensitivity analysis feature does permit you to consider the effects of different levels of load growth.

Load Duration Curves

The basic quantity upon which demand parameters are based is an hourly load, which measures the average level of demand over one hour. Hourly loads are given in megawatts (MW). Alternatively, load data can be input at 5-minute intervals. For a long-term planning model such as EGEAS, however, it is prohibitively time-consuming to perform all the analysis on an hourly basis. Therefore, the system demand for each period of time (year or subperiod) is represented by

three parameters - peak load, energy and a load duration curve. The peak load, expressed in MW, is the maximum of all the hourly loads within the period. The energy, in megawatt hours (MWh), is the sum of all the hourly loads; this sum is usually divided by 1,000 to convert the energy into gigawatt hours (GWh). The load duration curve (LDC) is the complement of the cumulative probability distribution for the hourly loads. For a given load level, the LDC gives the probability (or fraction of time) that the system demand exceeds that value.

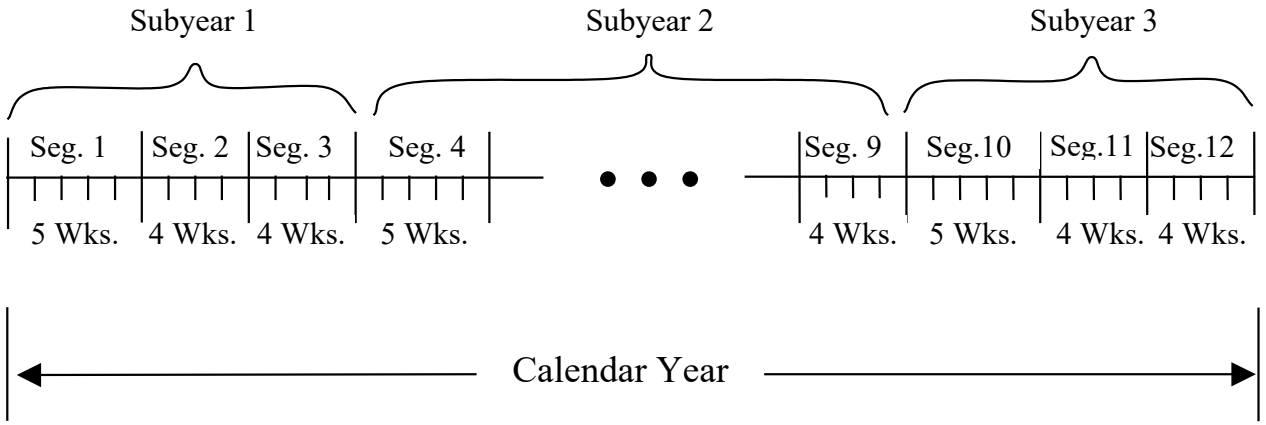

a. Segments and Subyears

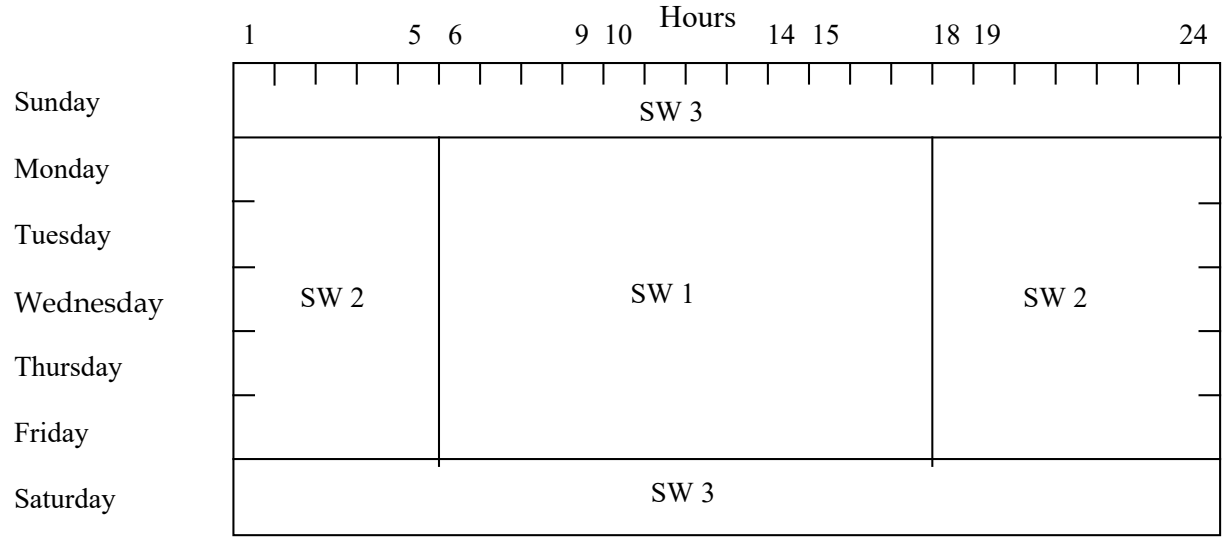

b. Subweeks

Figure 2-1
Subperiod Definition Example

Figure 2-2 illustrates the conversion of a time-dependent curve to a load duration curve. The left half of the figure shows a typical daily variation in power demand. Key parameters associated with the load distribution are the peak (maximum) and minimum load (the highest and lowest points on the curve) and the total energy (the area under the curve). To obtain the LDC shown in the right half, the loads are first re-ordered from the highest to the lowest and then the time axis is converted to a percent of time or probability. Note that the re-ordering does not change the peak or minimum load values, and the energy can be obtained from the LDC by multiplying the area by the number of hours in the time period.

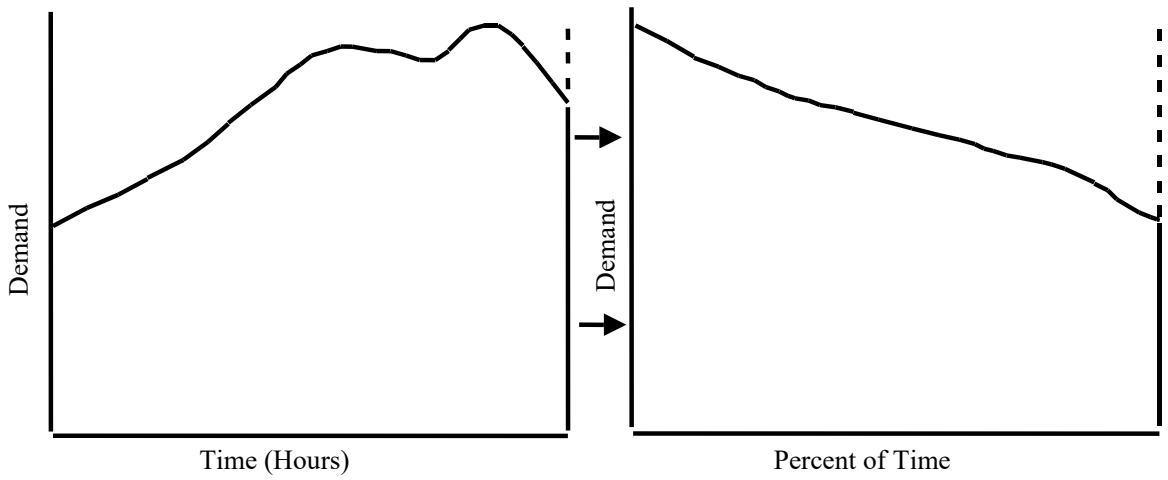

a. Time-Dependent Load Curve for a Typical Day

b. Load Duration Curve

Figure 2-2
Conversion of Time-Dependent Curve to Load Duration Curve

One other parameter associated with a load duration curve is the load factor. The load factor is simply the ratio of the actual energy for a period to the energy required if the load were constant at the peak load. The load factor is also the area under the LDC when the demand and time axes are both normalized to 1.0. The load factor can be calculated using the formula:

$$\text{LOAD FACTOR} = \frac{\text{ENERGY}}{\text{PEAKLOAD} * \text{HOURS}}$$

Figure 2-3 shows how conventional generating units can be operated to meet the system demand. The least expensive units to operate (e.g., nuclear) are loaded first and operate all the time. The most expensive units (e.g., natural gas turbines) are loaded last and operate only for a small portion of the time. Note that the conversion from a time-dependent curve to an LDC does not change how the units are loaded or how much energy they generate. Thus, the LDC can be used for production costing – determining the energy generated by each unit, and the expected amount of energy that cannot be served. However, because the time sequence of the hourly loads is lost, EGEAS does not model ramp rates, start-up costs, and minimum down-time; and EGEAS loses some accuracy in modeling storage units and interconnections. Production costing with load duration curves is described in Chapter 3.

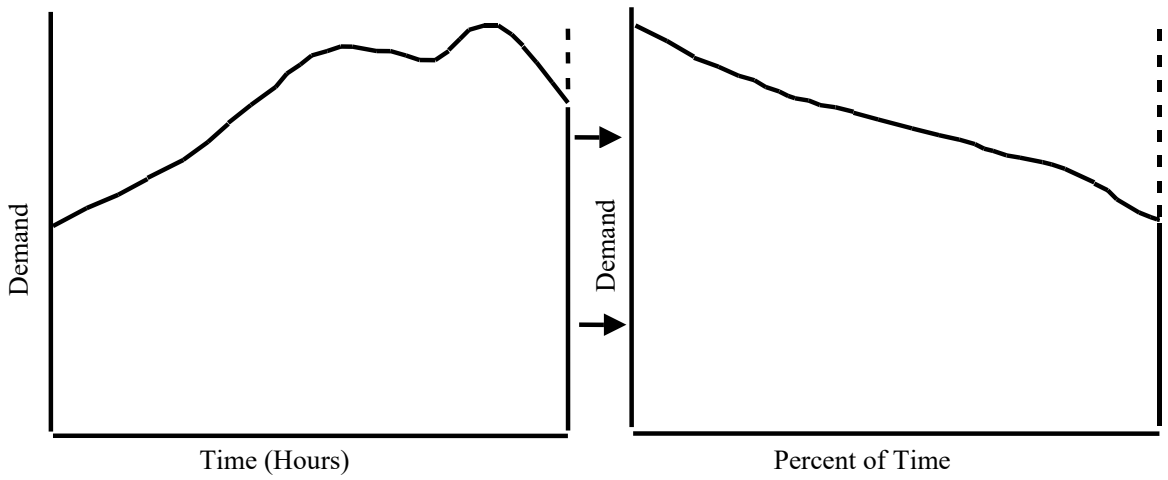

a. Time-Dependent Load Curve for a Typical Day

b. Load Duration Curve

Figure 2-3
Deterministic Operating Scheduling

Customer Classes

Customers, including residential, commercial, and industrial, all place demands for electricity on their electric utility. These individual demands can be represented on an individual basis or summed to obtain the total customer load as shown in Figure 2-4. Up to five different customer classes (also called load components) can be represented separately in EGEAS. Transmission and distribution losses should be included in the load for each class.

There are two primary reasons for representing the customer classes separately. The first reason is to estimate the revenues associated with each class. When a DSM program is added for a particular class, EGEAS can account for the loss in revenues for that class. The second reason is to model price elasticity. When modeling elasticity, EGEAS changes the system load based on increases in the cost of electricity relative to inflation. Price elasticity measures vary by customer class, so accurate representation of these changes in load necessitates working with individual class loads.

Interconnected Systems

For multi-area analysis, you must supply load data for each area or system. EGEAS allows a maximum of 10 areas, labeled as systems A through J, where system A is the primary system being analyzed. Only system A may have multiple customer classes. Also, system A is the only system that for which you can specify purchase and sale contracts or non-dispatchable technology resources.

EGEAS uses the loads for all the interconnected systems (including system A) to produce load duration curves for the combined system (pool). For information on how EGEAS constructs the pool LDC, see the Methodology section below. For information on multi-area modeling in EGEAS, refer to Chapter 8.

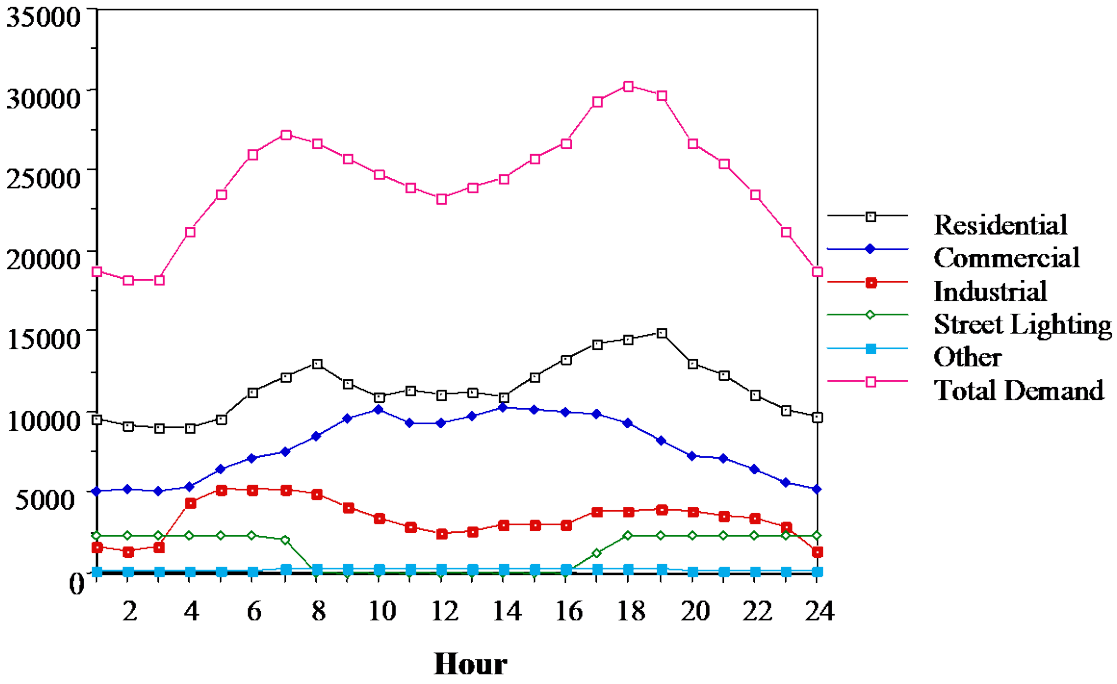

Figure 2-4
System and Class Load Curves

Input Requirements

You must supply hourly loads for every hour of the year. Data can be supplied in standard EEI format, an extended format allowing for input of larger values, or a one-column format with just one value per record. There is also an option for providing a file with data at 5-minute intervals. If desired, you can use the same set of hourly loads for multiple years. You may also specify peak loads and/or annual energies and EGEAS will scale the hourly loads to fit those parameters. If you do not specify a new peak load or energy, EGEAS will use the hourly loads as input. Actually, there are two ways to input peak loads and energies. While there is some redundancy in allowing both methods, there are some differences in how the loads are processed, as described in the methodology section below (steps 2 and 7).

Methodology

In EGEAS, the system load is used for production costing and reliability calculations. With features such as individual customer classes, purchase and sale contracts, non-dispatchable technologies, multi-area modeling, sensitivity analysis, and price elasticity, the load can get modified in different ways at numerous places throughout the analysis. The following outline shows the major ways in which the load for a particular year is manipulated and used. Many of the steps are optional, but when included, the operations are performed in the order shown.

1. Determine the annual peak load and energy from the hourly loads. The peak load is simply the maximum hourly value, and the energy is the sum of all the hourly values. When there are multiple customer classes, this step is performed for each class. When there are multiple areas, the step is also performed for each of the interconnected systems (B to J).

2. Modify the hourly loads to fit a new peak load and/or energy load. When there are multiple customer classes, this process can be performed for each of the classes. When there are two or more areas, the loads can be modified for the interconnected systems as well. To illustrate how the loads are modified, let

   P_0 = peak load from hourly loads (MW)

   E_0 = annual energy from hourly loads (MWH)

   P_1 = peak load supplied by user (MW)

   E_1 = annual energy supplied by user (MWH)

   H = hours in a year (hours)

   L = hourly load

   L' = modified hourly load

   If only a peak load value has been input, all the hourly loads are multiplied by the ratio of the input peak load to the peak load calculated from the hourly loads:

   $$L' = (P_1 / P_0) * L$$

   If only an energy value has been input, all the hourly loads are multiplied by the ratio of the input energy to the energy from the hourly loads:

   $$L' = (E_1 / E_0) * L$$

   If both the peak load and energy values have been input, then a linear transformation is applied to all the hourly loads such that the new peak load and energy will agree with the input values:

   $$L' = a * L + b$$

   where:

   $$a = \frac{E_1 - H * P_1}{E_0 - H * P_0} \qquad b = \frac{P_1 * E_0 - P_0 * E_1}{E_0 - H * P_0}$$

   If neither a peak load nor an energy value is input, this step is bypassed.

1. Combine the customer classes to get the total System A load. Loads are added together at the hourly level. This step is performed only if multiple customer classes have been defined.
2. Modify the total System A hourly loads to account for purchase and sale contracts. This step is performed only if purchase and/or sale contracts have been specified. For information on this process, refer to the Purchase and Sale Contracts section of Chapter 3.
3. For each defined subyearly time period, determine the load duration curve (LDC) from the hourly loads. The LDC for a given period is based on the distribution of the hourly values for that period. This process is performed for the total System A load, and for as many interconnected system loads (B to J) as have been specified.

4. For each defined subyearly time period, determine the joint probability distribution for the following hourly variables:

   System A total load (after purchase and sale contracts)

   Customer class loads (before purchases and sales)

   Interconnected system loads (B to J)

   Non-dispatchable technology per-unit output

   The total System A load variable is always present. The customer class loads are included when a model user specifies a load modification option; users should specify this option when they intend further changes to the load (see Step 7 below). The interconnected system loads and NDT variables are present when those features are modeled. This step is bypassed if the System A load is the only variable. From the joint probability distributions, EGEAS determines the parameters needed for modifying the LDC to account for NDT contributions, and later changes to the load itself. EGEAS also determines the parameters needed for constructing the LDC for the pool load, incorporating the loads from all the areas. For more information, refer to the Non-Dispatchable Technology section of Chapter 3.

5. Modify the annual peak load and energy. Note there are two places where the annual peak load and energy can be modified: (i) in Step 2 above; and, (ii) in Step 7 here. However, the primary purpose and the methodologies used for these modifications are different. At Step 2, the focus is on scaling historical load shapes to match forecasted peak loads and energies; changes are made to the hourly load values. At Step 7, the focus is on investigating different load forecasts; changes are made to the load duration curves. These changes may result from user input, from sensitivity analysis on load growth, or from modeling the effect of price elasticity. User input and sensitivity analysis may be applied to individual customer classes for System A or to the total system load for any of the interconnected systems (B to J); price elasticity applies only to the customer class loads for System A. In a given analysis, EGEAS may consider many different levels of load, so it is not practical to modify the loads at the hourly level. Thus, Step 7 deals with load duration curves. There are three possible sources for a new peak load or energy:

   User input. If the peak and energy are to be modified with sensitivity analysis (see below), EGEAS users must specify a peak load and energy at this point, even if one wants to keep the values supplied earlier. If a user wishes, they may specify a peak load and energy here even if they do not plan to use sensitivity analysis.

   Sensitivity analysis. Users may specify many different values for the peak load and energy for EGEAS to use in various combinations. This feature is useful for investigating the effects of different load projections on optimal expansion plan costs.

   Price elasticity. If an EGEAS user models price elasticity, EGEAS may need to change the load in later years of the analysis to reflect changes in the system average rate in the earlier years. When EGEAS is determining optimal expansion plans or simulating multiple plans, different load levels may result for different plans.

6. Modify the load to fit the new peak load or energy from Step 7. EGEAS uses several different methods to change the load, depending on what values are specified, and what modeling features have been selected. The various methods are as follows:

Only a new peak load is specified. EGEAS scales the load to the specified annual peak. If multiple subperiods have been defined, the peak load in each subperiod is increased or decreased by the same percentage as the annual peak. The annual and subperiod energy values all change by the same percentage as the peak loads. The actual load shape in each subperiod stays the same.

Only a new energy is specified. EGEAS scales the load to the specified energy. As with the peak load option above, the subperiod peak loads and energies all change by the same percentage and the individual load shapes stay the same.

Both a peak load and energy are specified, and there are no subperiods, multiple customer classes, purchase or sale contracts, non-dispatchable technologies, or multiple areas (systems). A linear transformation is applied to the annual LDC to obtain the desired peak and energy. The transformation is similar to that applied in Step 2 when trying to match both peak and energy.

Both a peak load and energy are specified, and there are subperiods, multiple customer classes, purchase or sale contracts, non-dispatchable technologies, or multiple areas. For any given segment of the year, the loads are scaled up or down, but the pattern of loads within the segment is not changed. The procedure is as follows: First, scale the loads to the desired annual peak load. The peak load and energy in each subperiod are increased or decreased by the same percentage as the annual peak load. At this point, the load is fixed for every segment whose peak load matches the annual peak load. Second, the loads in the remaining segments (those whose peak is less than the annual peak load) are scaled up or down to achieve the desired annual energy. The amount by which the loads in a particular segment are scaled is proportional to the difference between the initial peak load for that segment and the initial annual peak load. Segments with the lowest peak load are scaled more than segments whose peak is close to the annual peak load.

As an example, assume there are 4 segments of 13 weeks each and assume the starting peak loads (MW) and energies (GWH) are as follows:

Segment	Peak Load (MW)	Energy (GWH)
1	840	1,120
2	980	1,370
3	1,000	1,530
4	960	1,220
	Total	**5,240**

The total annual energy is 5,240 GWH. Next assume that the desired new peak load is 1,100 MW and the desired new energy is 6,000 GWH. First, EGEAS scales the loads to the new peak load:

Segment	Peak Load (MW)	Energy (GWH)
1	924	1,232
2	1,078	1,507
3	1,100	1,683
4	1,056	1,342
	Total	**5,764**

The resulting total annual energy is 5,764 GWH. Second, EGEAS scales the load in segments 1, 2, and 4 to obtain the desired annual energy without changing the annual peak:

Segment	Peak Load (MW)	Energy (GWH)
1	1,053	1,404
2	1,094	1,530
3	1,100	1,683
4	1,088	1,383
	Total	**6,000**

Note that the peak for segment 1 was increased by 129 MW, approximately 8 times the increase of 16 MW for segment 2, because the starting peak for segment 1 (924 MW) was about 8 times further from the annual peak than the starting peak for segment 2 (1,078 MW). The final segment energies sum to the desired annual value of 6,000 GWH.

In some cases, EGEAS is not able to modify the load correctly. If the specified new peak load and energy result in an annual load factor which is very different from the initial load factor, it may not be possible to scale the segment loads sufficiently to achieve the desired result. (EGEAS prints an error message to inform users of the problem). In such a case, users may need to change the load at an earlier stage (e.g., Step 2), or modify the hourly loads external to EGEAS.

7. Modify the System A load for non-dispatchable technology contributions. The load is also modified if a new peak load and/or energy has been specified for one or more customer class loads (in Step 7). The LDC in each subperiod is modified to account for the contributions of NDT units that have been installed and are available in the current

year. The load modification process uses the parameters derived from the joint probability distributions that were determined in Step 6. This step is performed only if NDTs are modeled, or if the customer class loads have been modified. For more information, refer to the Non-Dispatchable Technology section of Chapter 3.

8. Determine the LDC for the pool load. This step is performed whenever multi-area modeling is performed. For more information, refer to Chapter 8.
9. Perform production costing and reliability calculations. Using the net load after NDTs (if any), EGEAS uses production costing to determine the expected generation of the conventional generating units (and certain DSM resources). Production costing also yields the expected unserved energy. For multi-area analysis, production costing is performed for each of the interconnected systems (B to J), and for the pool, as well as for System A. Loss of load probability calculations are also based on the net load after NDTs. The reserve margin calculation uses the peak load before NDTs, but the NDT contributions are accounted for. For more information on reliability, refer to the section on Reliability Measures. For more information on production costing, refer to Chapter 3.

Resources: Supply-Side and Demand-Side

EGEAS is designed to model supply-side and demand-side resources on an equal footing for integrated resource planning. Supply-side technologies represent generating units that supply the electricity to meet the system load. Demand-side technologies can be designed to increase or decrease the load, or to shift some of the load to different hours. Supply-side and demand-side resources share many of the same characteristics. The subsections that follow describe these characteristics and also indicate where the supply-side and demand-side resources differ. The following section on Other Resources describes the contributions from purchase and sale contracts and from interconnections with other utilities or power pools.

Resource Type and Status

Electricity to supply the system load is provided by four types of generating units: thermal, hydro, storage, and non-dispatchable technology (NDT). There are also four equivalent types of demand-side management (DSM) resource: dispatchable, energy limited, storage, and NDT. On the supply side, modeling is performed with individual generating units rather than with entire power plants or stations. On the demand side, a resource can represent a single DSM program or several programs grouped together by program marketing applications, by resource life, by customer class segment, or by other user-defined grouping.

Thermal units are those which can be dispatched as needed, up to their full capabilities, to meet or reduce the system load. As with all types of resources, however, these units are subject to mechanical outage conditions and scheduled maintenance requirements. Examples include nuclear, fossil (i.e., coal, oil, and gas), combined cycle, and combustion turbine technologies. For DSM, the dispatchable type can be used to represent conservation or strategic marketing programs that have load shapes that are constant over a daily period.

Hydro units, also known as limited energy plants, are dispatched in the same manner as thermal units, but the amount of energy generated is limited by conditions external to the plant. One example is a conventional (reservoir) hydroelectric plant, whose output is limited by the reservoir size and stream flow. Another example is a fossil unit whose output is limited by the available fuel supply or by emission constraints. For DSM, the energy limited type can be used

for programs that result in temporary curtailment of electric usage without any associated load rebound. Examples include interruptible rates, stand-by power generation, and other peak clipping programs.

Storage units take energy from the electric grid, usually during off-peak hours, and return energy when more capacity is needed, usually during peak load hours. Examples include pumped hydro, batteries, compressed air, and superconducting magnets. For DSM, the storage type can be used for technologies that add or save energy in one period, and require some energy payback in another period. Examples include commercial cool storage, heat storage, and other load shifting programs.

Non-dispatchable technology (NDT) units are characterized by (i) generation, or changes in demand, (ii) dependent on the time of day or correlated with the system load, and (iii) units cannot be dispatched as desired. Examples include solar, wind, run-of-the river hydro, and cogeneration. NDT resources require hourly data to be entered on a per-unit basis. Both positive and negative values are accepted. For DSM, the NDT type can be used for load control programs to capture the effects of load rebound immediately after the control period. Examples include water heater and air conditioner cycling control programs. Also, the NDT type can be used for load reduction programs based on efficiency improvements or behavioral changes induced by rates. Examples include high-efficiency appliances, heat pumps, insulation, and time-of-use rates.

All resources in EGEAS are classified by their status as well as by their type. Existing and committed status are for units or technologies whose installation dates are given by user input; existing units are installed before or in the base year, while committed units are installed after the base year. Generic status is for resources that may be installed to meet increased system demand, replace retiring units, or even reduce demand. The number of generic units chosen and their installation dates are determined during the optimization process, but constraints may be placed on the number of units available or the years in which resources can be installed.

All resources are assumed to be installed on January 1 of the installation year, and to retire on December 31 of the retirement year. For example, a unit installed in 1966 with an operating life of 40 years would begin generating energy on January 1, 1966 and stop generating on December 31, 2005.

Capacity

Four different capacities are associated with each generating unit or DSM resource; each is measured in megawatts (MW), and is assumed to represent the net power available to meet system demand, or to reduce demand in the case of DSM. Rated capacity is the capacity for which the plant or DSM program was designed.

Rated capacity is used for calculating capacity factors, capital costs and fixed O&M costs; and for describing plants on output reports. For resources, such as DSM, that are installed in stages over several years, the rated capacity can be changed from year to year. When the rated capacity increases, the increase is treated as new installed capacity with a new stream of capital fixed charges. (Refer to the Expansion Plan Costs section for more details).

Operating capacity measures the net power or megawatts available at a particular time. The operating capacity can vary both yearly and seasonally; however, seasonal variation is not

allowed for non-dispatchable technology units. The operating capacity is used in determining the energy generated by a unit whenever production costing is performed.

Emergency capacity is used to calculate the system loss of load probability. The emergency capacity represents the limit to which a unit can be pushed, or program marketed when an emergency situation arises. The value input will depend on whatever assumptions a user wants to make in determining loss of load probability. The emergency capacity can be less than or greater than the rated capacity. Annual and seasonal variations in the emergency capacity must coincide with those for the operating capacity, so the ratio of operating and emergency capacities remains constant.

Reserve capacity is used to calculate the system reserve margin. Yearly variations in the reserve capacity are allowed, and need not coincide with variations in the operating and emergency capacities. Seasonal variations must follow the same pattern as those for the operating capacity. The use of four capacities provides model users with considerable flexibility. However, it should be noted that it is often appropriate to assign the same constant value to all four capacities.

Two additional factors affect the capacities used by EGEAS. The first is the percentage of the resource owned by the electric company being modeled. When not specified, this parameter is assumed to be 100 percent. If a resource is not fully owned, all of its capacities are lowered in accordance with the ownership percentage. For example, a 400 MW unit that is 60 percent owned is treated as if it were a 240 MW unit. Heat rates, outages and other characteristics are not affected. The second factor is applicable only to NDT type resources. With plants of this type, such as solar and wind generators, there are often many small units rather than one or two large units. The capabilities of EGEAS would be strained or exceeded in trying to handle a large number of such units individually. Therefore, a resource can represent a group of units. The number of individual identical units is specified and the capacity is given for one individual unit. EGEAS treats the units as independent entities for modeling, but groups them for reporting purposes.

For increased modeling accuracy, a unit can be subdivided into two or more loading blocks. A maximum of five blocks per unit is allowed. Each loading block represents a fraction of the overall unit capacity. A unit with N blocks has N+1 operating states because the unit could be entirely out of service. For example, a 1000 MW unit divided into four loading blocks of 500, 200, 200, and 100 MW could operate at capacities of 0, 500, 700, 900, and 1000 MW. Block capacities are specified as fractions of the total unit capacity (e.g., 0.5, 0.2, 0.2, and 0.1). The same fractional breakdown applies to rated, operating, reserve and emergency capacities. The use of loading blocks permits more accurate fuel cost calculations based on incremental heat rates for each block, and more accurate reliability calculations based on partial unit forced outage rates. There is a restriction that storage and NDT units cannot be divided into loading blocks.

Heat Rate

The fuel requirements for a generating unit are determined by using a heat rate, which measures the fuel energy required per unit of net energy generated. Heat rates are required for thermal units, but are optional for hydro and storage units. They are not used for NDT units. Heat rates are given in BTU/kWh or equivalently MBTU/GWH. Yearly changes in the heat rate are allowed, but seasonal variations are not.

When a unit is loaded as a single block, EGEAS uses the net full load heat rate, which gives the BTU/kWh for the unit operating at full capacity. When loading blocks are modeled, separate heat rates can be specified for each block. For the first block, EGEAS uses an average heat rate, which gives the BTU/kWh for the unit operating at the first block capacity. For a higher block, EGEAS uses an average incremental heat rate, which gives the average BTU needed to produce an additional kWh of generation within the given block.

As an example, consider a 600 MW unit with the following characteristics:

Output (MW)	Input (MBTU/HR)
0	506
100	1,352
230	2,492
350	3,585
450	4,526
600	5,990

The full load heat rate is:

$$\frac{5{,}990 * 1000}{600} = 9983\,\text{BTU/kWh}$$

(where the 1000 converts from MBTU/MWH to BTU/kWh).

If the unit is divided into five blocks with capacities of 100, 130, 120, 100, and 150 MW, then the average heat rate for the first block is:

$$\text{Block 1}: \frac{1352 * 1000}{100} = 13{,}520\,\text{BTU/kWh}$$

The average incremental heat rates for the higher blocks are:

$$\text{Block 2}: \frac{2492 - 1352 * 1000}{230 - 100} = 8{,}769\,\text{BTU/kWh}$$

$$\text{Block 3}: \frac{3585 - 2492 * 1000}{350 - 230} = 9{,}108\,\text{BTU/kWh}$$

$$\text{Block 4}: \frac{4526 - 3585 * 1000}{450 - 350} = 9{,}410\,\text{BTU/kWh}$$

$$\text{Block 5}: \frac{5990 - 4526 * 1000}{600 - 450} = 9{,}760\,\text{BTU/kWh}$$

The input/output curve based on this data is shown in Figure 2-5.

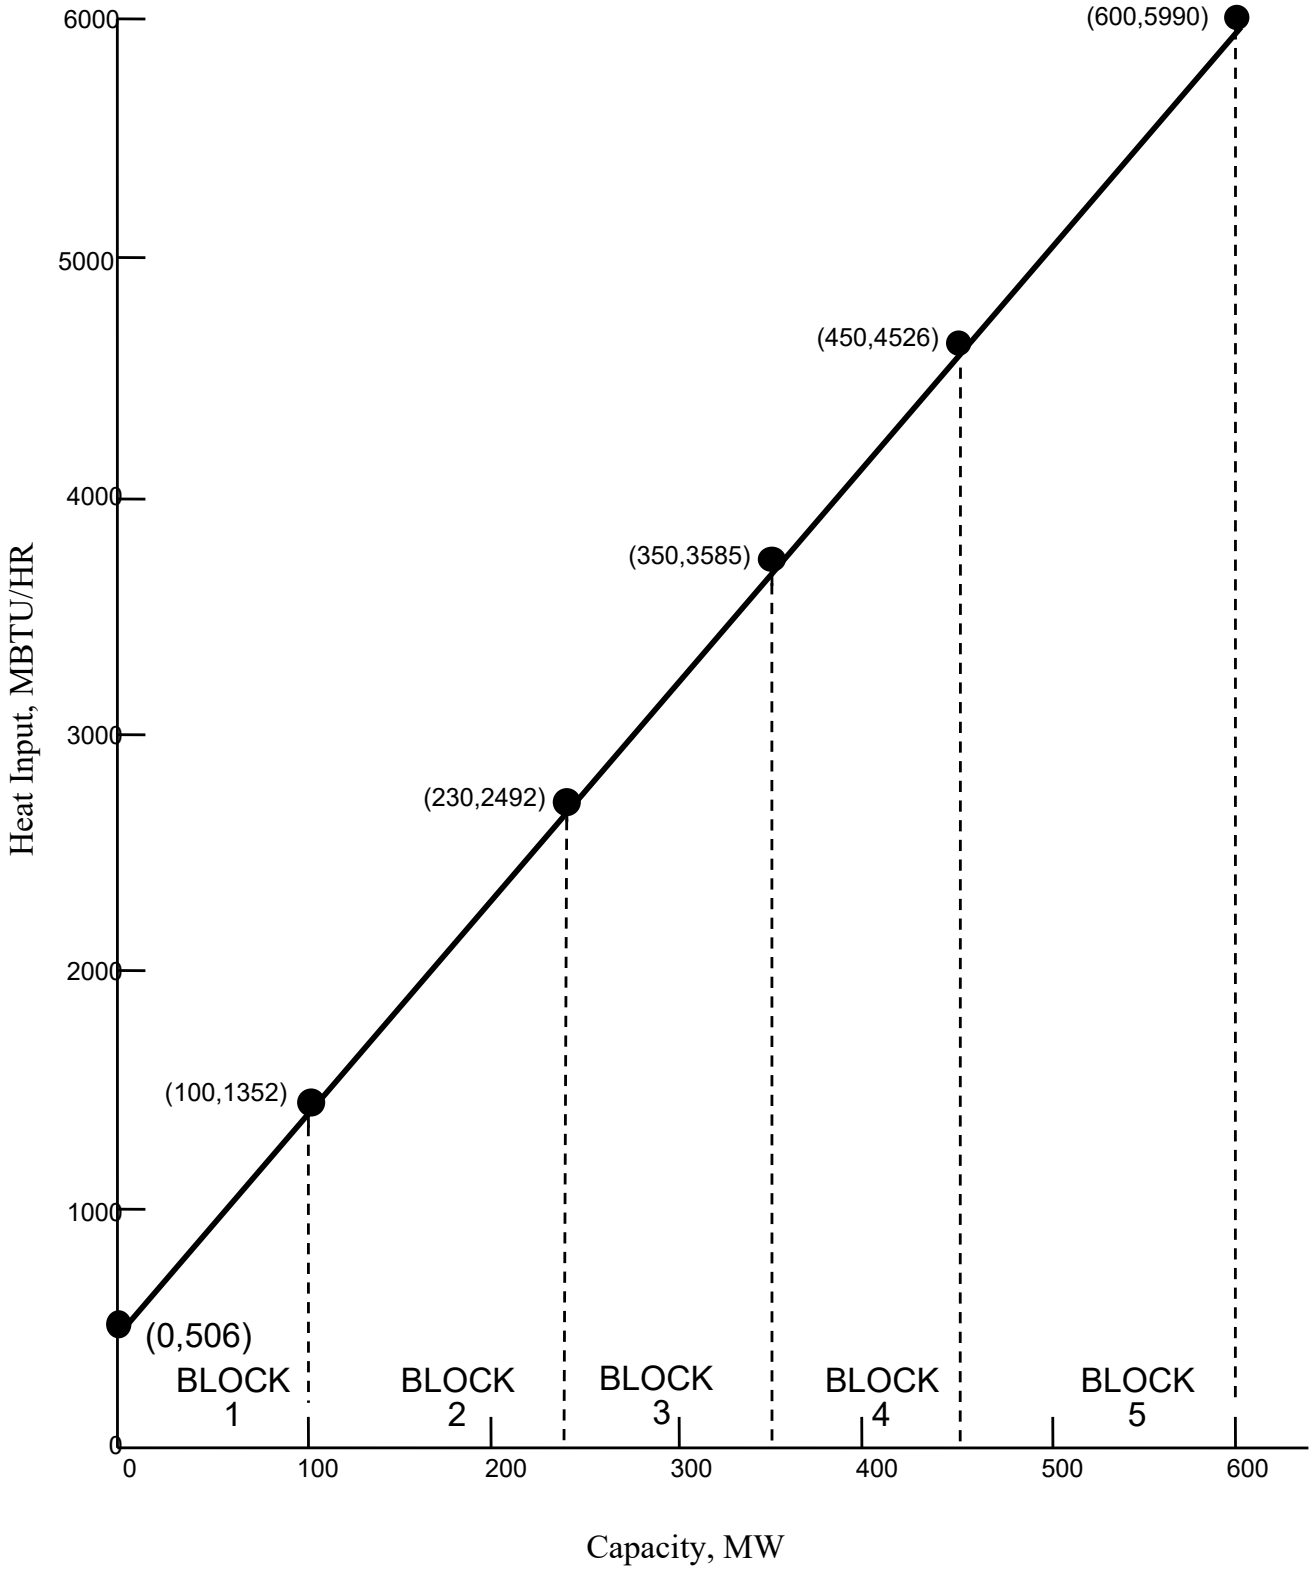

Figure 2-5
Heat Input-Output Curve

In the output reports, the average heat rate of a unit is given. If loading blocks have been modeled, this average is determined from the incremental heat rates of the individual blocks using the energy generated by the blocks as weighting factors. If loading blocks have not been modeled, the average heat rate is assumed to be the full load heat rate. For the example given above, assume that the blocks generate the following amounts of energy:

Block	Energy (GWH)
1	700
2	900
3	810
4	670
5	990
Total	**4,070**

Assuming that loading blocks have been modeled, the average heat rate is:

$$\frac{(700 * 13{,}520) + (900 * 8{,}769) + (810 * 9{,}108) + (670 * 9{,}410) + (990 * 9{,}760)}{4{,}070} = 10{,}000$$

If loading blocks were not modeled, the average heat rate would be taken as 9,983, the full load value.

Forced Outage Rate

Associated with each generating unit is a forced outage rate, the fraction of time the unit is out of service due to mechanical failure. The rate also can be used to determine the amount of time DSM technologies will be available to either increase or decrease customers' use of electricity. Scheduled maintenance requirements are handled separately. The complement of the forced outage rate is availability, i.e., the fraction of time a unit is operational. The amount of energy a unit can generate is limited by the product of its operating capacity, its availability, and the number of hours in the time period. Often the forced outage rate is higher in the first few years of a unit operates. Thus, the forced outage rate is allowed to vary on a yearly basis to account for these immature unit characteristics. The outage rate is assumed to remain constant over the course of a given year.

For more detailed simulation, thermal and hydro units can be divided into loading blocks. Separate outage rates can be assigned to individual blocks, permitting the modeling of partial outage states. When one block is forced out, then all higher blocks cannot operate. Thus, if a unit is divided into N blocks, there are N+1 possible operating states, including the state where all blocks are available. The following table illustrates the possible states for a unit with 4 blocks:

State	Block 1	Block 2	Block 3	Block4
1	X	X	X	X
2	Available	X	X	X
3	Available	Available	X	X
4	Available	Available	Available	X
5	Available	Available	Available	Available

Because a unit must always be in one of these states, the probabilities assigned to the states must sum to 1.0. In EGEAS, a partial outage is represented by a block outage rate, which is defined as the probability that a given block is forced out, and that all prior blocks are available. In other words, the block outage rate is the probability that the given block is the lowest block that is forced out. The block outage rate for block i is simply the probability of state i.

The block availability is defined as the probability that a given block is operational. Because a block cannot operate if any lower blocks are forced out, the block availability is one minus the sum of the block outage rates taken over all blocks up to the given block. For the first block of a unit, the block availability is always the complement of the block outage rate.

The equivalent forced outage rate (EFOR) for the entire unit is the mean fraction of total unit capacity forced out. EFOR is calculated as the weighted average, over all the states, of the capacity forced out, divided by the total capacity of the unit. If loading blocks are not modeled, then the equivalent forced outage rate is used as the forced outage rate of the unit. If a model user wants to use loading blocks to more accurately model heat rates, but wants to treat the unit as a whole for outages (the unit is either available with full capacity or not at all), then the block outage rate for the first block should be set to the unit's forced outage rate, and the block outage rates of higher blocks should be set to zero.

As an example, consider a 1,000 MW unit with the following states:

State	Available Capacity (MW)	Capacity Forced Out (MW)	Probability
1	0	1000	0.05
2	500	500	0.10
3	700	300	0.20
4	900	100	0.40
5	1000	0	0.25
		Total	**1.00**

The equivalent forced outage rate for this unit is:

$$\frac{\text{EFOR} = [(.05)*1000] + [(.10)*500] + [(.20)*300] + [(.40)*100] + [(.25)*0]}{1000} = 0.2$$

In the EFOR calculation, the numbers shown in parentheses are the probabilities for the states; the first four probabilities are the block outage rates. The last term, corresponding to the highest state, can be omitted because no capacity is forced out.

To model this five state unit in EGEAS, you would specify a forced outage rate of 0.2 for the unit as a whole and define four loading blocks with the following characteristics:

Block	Capacity	Block Outage Rate	Block Availability
1	500	.05	0.95
2	200	.10	0.85
3	200	.20	0.65
4	100	.40	0.25

The block information can be obtained from the state table above. The capacity for block i is the difference in available capacity between state i and state i+1 (e.g., for block 2: 700 - 500 = 200). The block outage rate for block i is simply the probability of state i (e.g., for block 2: 0.10). For the first block, the block availability is the complement of the block outage rate (1.0 - 0.05 = 0.95). For higher blocks, the block availability is the availability for the previous block minus the outage rate for the current block (e.g. for block 2: 0.95 - 0.10 = 0.85). Block availabilities are calculated internally by EGEAS.

Maintenance

The term maintenance is used in EGEAS to refer to planned periods when a unit is to be taken out of service. This is in contrast to forced outages, which may be anticipated, but are not planned. Maintenance requirements are always given as integral numbers of weeks per year. For a given unit, the number of weeks required for maintenance may vary from year to year. When analysis is performed on an annual basis, the capacity and cumulants of a unit are derated by the fraction of time that the unit is on maintenance. When analysis is performed on a subperiod basis, the maintenance is allocated to one or more segments or subyears. EGEAS supports up to two maintenance periods per year. If one period is specified, you can specify when it starts, or let the EGEAS maintenance scheduler choose. If two periods are specified, you must specify the starting point for both periods.

Fuel

For every thermal unit, users must specify one or two fuel types. For hydro and storage units, users have the option to specify a single fuel type. Users may define general fuel types that are used by many units, or separate fuel types for each generating unit. For each fuel type, users specify a unit of mass (e.g., tons of coal or barrels of oil), and a heat content (in millions of BTU per unit of mass). Users also specify the amount of fuel available per year (in units of mass) and the fuel cost (in dollars per million Btu). The available fuel and its cost can be changed annually and seasonally.

For a unit with a single fuel type, the fuel usage is given by the equation:

$$\text{FUEL USAGE} = \frac{\text{ENERGY} * \text{HEAT RATE}}{\text{HEAT CONTENT}}$$

For a unit with two fuels, primary and secondary, you can specify a target percentage for the secondary fuel. In the absence of fuel constraints, EGEAS assumes that the target percentage of a unit's generation is fueled with the secondary fuel, with the remaining generation provided by the primary fuel. If you limit the amount of fuel that is available, EGEAS can change the fuel percentage based on the available supply of both fuels. You may, however, specify minimum and maximum limits on the second fuel percentage.

Limited Energy and Storage Plants

The amount of energy a unit can generate is limited by the product of its capacity, its availability, and the number of hours in the time period. In addition, for hydro and storage units, users can specify an energy limitation, which may be more restrictive than the capacity-availability constraint. For a hydro unit, this energy constraint can be used to represent the limitations imposed by the amount of flow in the river and the reservoir capacity. For a storage unit, the energy constraint can represent the storage capacity of the unit. When there are multiple charging/discharging cycles within a time period, the energy limit should be the product of the number of cycles and the storage capacity per cycle. In every case, the limitation must be expressed in GWH. For a given unit, the energy limit may vary from year to year, and when subperiod analysis is performed, the energy may be distributed throughout the year as needed to model seasonal variations. For hydro units, the energy limit can be allocated to different subweeks. For storage units, charging and generating operations can be allocated to different subweeks. For example, charging could be performed at night, and generation during the daytime.

Two additional parameters apply only to storage units. The first parameter is storage efficiency, which indicates the ratio of the energy generated to the energy stored. Because no storage unit is completely efficient, the energy generated is usually less than the energy stored and the efficiency is less than 100 percent. However, compressed air energy storage units have storage efficiencies greater than 100 percent because they burn fuel in addition to taking charging energy from the system. For a given unit, the storage efficiency is assumed to be time invariant. The energy limitation is always applied to the generating side, as opposed to the charging side, of the unit. The second parameter is charging capacity, which measures the capacity at which energy can be stored. The operating capacity of the unit is used for generation. The ratio of the charging capacity to the operating capacity is assumed to remain constant so both capacities follow the same patterns of annual and seasonal variation. The forced outage rates and maintenance schedules for the charging and generating sides of the unit are assumed to be the same.

Non-Dispatchable Technology Resources

The primary distinguishing characteristic of NDT units or resources is that their generation or reduction in load is dependent either on time of day or the customer load. Modeling this dependence accurately requires that EGEAS develop a joint probability distribution for the customer load and all non-dispatchable technologies considered. Therefore, whenever NDT units are modeled, users must supply hourly generation parameters for each technology. These parameters, called utilization levels, provide the generation per unit of capacity for each hour of

the year. Utilization level values assume that a unit is always available from a mechanical standpoint; forced outages and maintenance requirements are incorporated later in the analysis. Utilization levels represent the effect of such parameters as solar radiation or wind speed. A distinction should be made between a non-dispatchable technology and an NDT unit. A technology is represented by a series of utilization levels, while an NDT unit includes a specific capacity, forced outage rate and costs. Many NDT units may share the same technology.

Expansion Plan Costs

For the cost of an expansion plan, EGEAS uses the present worth of revenue requirements, which include the fixed charges associated with capital investments, and the operating expenses associated with generation to meet the system load. For each year, these costs are summed, and then discounted to a base year to obtain a present worth cost. These present worth costs are summed over all the years of both the study and extension periods. For more detailed information on the handling of costs in the extension period, refer to the End Effects section later in this chapter.

Capital Fixed Charges

Capital costs, which represent the expense of building a new generating unit or capitalized DSM program costs, can be associated with existing, committed, and generic resources. If the assets of existing, or even committed, resources are already merged into an electric utility's rate base, users may omit their capital costs. However, users may want to include these costs for reporting purposes, or for inclusion in the calculation of feasibility measures such as the system average rate or the interest coverage ratio. For a generic resource, the decision to build a unit or implement a program depends heavily on the capital costs, which should therefore always be included. In addition to the initial costs, you may also specify on-going capital expenditures for resources that have already been installed.

You may choose between three methods for representing the capital costs: (i) levelized fixed charges; (ii) annual fixed charges; and, (iii) construction cost modeling. With the levelized fixed charges option, EGEAS calculates capital fixed charges by applying a levelized fixed charge rate to the installation cost in each year of the resources book life. With the annual fixed charges option, EGEAS calculates capital fixed charges for each year of the resource's book life, considering book and tax depreciation, interest rates, rates of return allowed on equity, and tax rates. With the construction cost option, EGEAS calculates the allowance for funds used during construction (AFUDC) and construction work in progress (CWIP) for each year of the construction period. Annual fixed charges are then calculated over the book life of the resource. For more information on how capital costs are calculated, refer to Chapter 7. For information on how EGEAS handles the end effects associated with capital costs, refer to the End Effects section later in this chapter.

For generic resources, you can specify two separate capital installation costs. Conceptually, units of a particular type are installed or implemented in groups, with one cost assigned to the first unit in each group, and the other cost assigned to all other units. Thus, if the group size is 4 units, the first cost is assigned to the first, fifth, ninth, etc. units and the second cost is assigned to all other units. Users determine the size of the group, as well as the two separate costs. Note that users can use just one cost for all units if desired. This feature can be helpful in modeling the extra costs, such as site preparation costs or initial marketing expenses, associated with the first unit of a

particular resource. For analyses with screening curves and Benders decomposition, which deal with fractional numbers of units, EGEAS uses a weighted average of the two costs for all units.

Capital costs for a resource are based on the rated capacity. In some cases, the rated capacity may change from year to year. For example, the size of a DSM program may increase over several years based on market penetration. When the rated capacity increases, the increase is treated as new installed capacity. Associated with the new capacity is a new stream of capital fixed charges. If the capacity decreases, there is no impact on capital cost. If the capacity decreases and then increases, EGEAS will include new capital costs only when the capacity exceeds the highest previous value. If user input results in capacity changes after the end of the study period, those changes are ignored. Below is an illustration:

Study Period: 1995 to 2004	
Unit Capacity	
Years	**Capacity (MW)**
1 to 3	100
4 to 5	50
6 to 8	80
9 to 15	150
16 to 30	200

For a unit installed in 1995, the installed capacity would be as follows:

Years	**Capacity (MW)**
1995 to 1997	100
1998 to 1999	50
2000 to 2002	80
2003 to 2004	150
Extension period	150

Capital costs would be determined for the following capacity additions:

1995	100 MW
2003	50 MW

Note that the increase from 50 to 80 MW in 2000 does not result in capital costs because the capacity is still below the 100 MW in 1995. Also, the capacity does not increase to 200 MW because such addition would come after the end of the study period.

Operating Costs

Generating unit operating costs and DSM option expensed-related costs consist of fixed operating and maintenance (O&M) costs, variable O&M costs, and fuel costs. These costs can be supplied for all resources - existing, committed, and generic.

Fixed O&M costs represent operating costs that are independent of the generation output of the unit or the savings attributed to the DSM program. These costs, which are expressed as dollars per kilowatt of rated capacity, are applied starting in the installation year of the resource and

continuing throughout the operating life of the resource. Fixed O&M costs are escalated from year to year. If the rated capacity changes from year to year, the fixed O&M cost in a particular year depends on the rated capacity value for that year. For situations where a resource (e.g., an energy sale), might produce revenues instead of costs, the fixed O&M cost can be negative.

Variable O&M costs represent a direct expense of generating energy or energy savings, and are expressed as dollars per MWH of generation or energy savings. In the case of storage resources, the variable O&M cost depends on the amount of energy generated; the cost is not applied to the amount of energy charged. For NDT resources, the variable O&M is applied to the net energy generated or net energy savings. Because the net energy for a DSM program can be negative, EGEAS allows the variable O&M cost to be positive or negative. A negative variable O&M cost can also be helpful when modeling an energy sale.

Fuel prices are specified in terms of dollars per million BTU ($/MBTU). To obtain the fuel cost for a resource, EGEAS uses the product of the fuel price, the resource's heat rate, and the energy generated or saved. For a storage resource, such as a compressed air energy storage unit, EGEAS uses the energy generated, not the energy charged. There are no fuel costs associated with NDT resources. When a thermal resource has two fuels, EGEAS determines the fuel cost based on the percentage of the unit's generation that is supplied with each fuel, and the individual heat rates and fuel prices. Fuel prices can be changed both annually and seasonally.

Detailed Costs

Detailed costs represent several additional costs and savings that can be associated with generating units or DSM programs. There are five categories: (i) DSM customer costs; (ii) rebound benefits associated with DSM programs; (iii) customer benefits from DSM technologies; (iv) transmission/distribution costs; and, (v) other costs. These costs and benefits can be specified in one of three ways: (i) as capital costs; (ii) as fixed O&M costs; or, (iii) as variable O&M costs. Different methods can be used for different categories of costs and for different resources. For a cost modeled as fixed or variable O&M, users can specify an expenditure pattern that describes how the costs change from year to year over the life of the unit. This is in addition to the trajectory that can be specified for all types of detailed costs to reflect price escalation.

Detailed costs are included in the total expansion plan cost. For optimization, detailed costs are included when the objective function is revenue requirements. Users can decide whether they should be included when the objective function is the levelized system average rate. Users also can specify whether different categories of detailed costs are discounted with the system discount rate or the customer discount rate.

Purchase Costs and Sale Revenues

For purchase contracts, users can specify capacity costs, in dollars per kilowatt of capacity per year, and energy costs, in dollars per megawatt hour of energy purchased. Similarly, for sale contracts, there are capacity revenues and energy revenues. For modeling flexibility, these costs/revenues can be either positive or negative; they can be escalated from year to year, and they may vary seasonally. When multi-area modeling is performed, a system may incur purchased energy costs or may receive revenues from energy sales. Refer to Chapter 8 for details on the available options for calculating the purchase and sale costs with multi-area modeling.

Unserved Energy Cost

No generating system is perfectly reliable, so there always is a possibility, however small, that the system demand cannot be fully met. EGEAS assigns a cost to this expected unserved energy, based on an input price ($/MWH). This price may represent the actual cost of emergency purchases from another system, or the cost to society of an electric supply outage.

Emission Costs

Depending on the emission types specified and emissions modeling options, EGEAS will calculate the cost of emissions. The cost may be based on all the emissions produced, or all the emissions above a certain limit. Users have many options for applying emission constraints and assessing emission costs. Refer to Chapter 9 for further information.

Reliability Measures

EGEAS uses four different measures of system reliability to determine how well the system can fulfill system demands for electricity:

1. Reserve Margin
2. Loss-of-Load Probability
3. Relative Reliability
4. Unserved Energy

Each of these measures is described below.

Reserve Margin

The simplest measure of system reliability is the reserve margin, which is the percentage by which the total generating capacity exceeds the annual peak load. Generating capacity is summed over all the units in the system (i.e., installed and not yet retired). EGEAS uses reserve capacities for the reserve margin calculation. As an example, a peak load of 10,000 MW and a generating capacity of 11,500 MW produce a reserve margin of:

$$\left[\frac{11{,}500 - 10{,}000}{10{,}000}\right] * 100 = 15 \text{ percent}$$

The reserve sharing option treats contributions from an interconnected system as an additional generating unit. For example, if 500 MW were available through reserve sharing, the reserve margin in the above example would increase to:

$$\left[\frac{(11{,}500 + 500) - 10{,}000}{10{,}000}\right] * 100 = 20 \text{ percent}$$

Purchase and sale contracts, described in Chapter 3, also may contribute to the reserve margin. For a purchase, the contract may be treated as an increase in supply or as a reduction in demand (i.e., shifting the load to another system). Continuing with the previous example, a 400 MW purchase treated as an additional supply would change the reserve margin to:

$$\left[\frac{(11{,}500 + 500 + 400) - 10{,}000}{10{,}000}\right] * 100 = 24 \text{ percent}$$

Treated as a reduction in load, the contract would change the reserve margin to:

$$\left[\frac{(11{,}500+500)-(10{,}000-400)}{(10{,}000-400)}\right]*100 = 25\ \text{percent}$$

Similarly, a sale contract can be represented as a decrease in system resources or an increase in system demand. For the above example, a 300 MW sale treated as a reduction in resources would change the reserve margin to:

$$\left[\frac{(11{,}500+500-300)-(10{,}000-400)}{(10{,}000-400)}\right]*100 = 21.88\ \text{percent}$$

Treated as an increase in demand, the sale would change the reserve margin to:

$$\left[\frac{(11{,}500+500)-(10{,}000-400+300)}{(10{,}000-400+300)}\right]*100 = 21.21\ \text{percent}$$

Demand-side management resources contribute towards reserve margin by reducing the peak load. For example, a system with a peak load of 10,000 MW, generating capacity of 11,500 MW, and DSM capacity of 400 MW would have a reserve margin of:

$$\frac{11{,}500-(10{,}000-400)}{(10{,}000-400)}*100 = 19.79\ \text{percent}$$

You can request that the reserve margin be calculated on an annual basis or on a subperiod (segment) basis. The annual calculation uses the annual peak load and the average reserve capacity for each resource. The subperiod calculation uses segment peak loads and segment values for reserve capacities, thus accounting for seasonal variations. In either case, reserve margin constraints are applied only on an annual basis, but the segment option, the annual reserve margin to be constrained is the minimum of the segment values.

Users also have an option to include or exclude maintenance in the reserve margin calculations. When the maintenance option is selected, the capacity for a resource is derated by the portion of time the resource is on maintenance in that period. When dealing with subperiods, EGEAS uses the maintenance schedule determined for each resource.

Loss of Load Probability

The second measure of reliability is loss of load probability (LOLP). Within EGEAS, LOLP is defined as the probability (fraction of time) that the system demand is not fully satisfied, or in other words, that the system demand exceeds the available capacity. When multiplied by the number of hours in a time period, the LOLP yields the loss of load hours. LOLP does not measure the magnitude of the loss. LOLP is calculated by convolving the unit outages onto the load duration curve (as described in Chapter 3). After all the units have been loaded, LOLP is the height of the equivalent load duration curve at the final capacity point (see Figure 3-4). For the LOLP calculation, EGEAS uses the emergency capacity for each resource.

Users can request that the LOLP be calculated for each subperiod within a year rather than just on an annual basis. Calculations for a particular period are based on the load and unit capacities for that period, thus accounting for seasonal variations. LOLP constraints are still applied on an annual basis, but the annual LOLP to be constrained is determined from the subperiod values. You also have an option to include or not include maintenance into the LOLP calculations. With the maintenance option, resource capacities are derated by the portion of time a unit is on

maintenance. When dealing with subperiods, EGEAS uses the maintenance schedule determined for each resource.

Relative Reliability

The third measure of reliability is based on the LOLP and is called relative reliability. To obtain the relative reliability, EGEAS compares the LOLP for a given year to the LOLP for a benchmark year. This benchmark year, often the first year of the study period or some year before the study period, is chosen by the model user. Users also can select a benchmark peak load which is used in the LOLP calculation. EGEAS cannot model all the different definitions of reliability used by electric companies, but the benchmark lets users set up conditions that exactly meet any company's reliability constraint. The relative reliability then indicates whether the system is more or less reliable than the benchmark in any given year. The actual definition of relative reliability is the ratio of the benchmark LOLP to the LOLP in the given study period year. A relative reliability greater than 1.0 indicates that the system is more reliable than the benchmark. As an example, a benchmark LOLP of 0.006 and a study year LOLP of 0.002 produce a relative reliability of 3.0.

Unserved Energy

The fourth measure of reliability is unserved, or unmet, energy. Unserved energy, expressed in GWH, represents the portion of the initial system energy that cannot be met by the system's generating units. The unserved energy is calculated during production costing, as described in Chapter 3. After all units have been loaded, the unserved energy is the area under the equivalent load duration curve (see Figure 3-4). For the unserved energy calculation, EGEAS uses the operating capacity for each resource.

End Effects

The artificial end of a planning period can distort analyses, and thus planning decisions based on those analyses. One way to minimize these end effects is to analyze a long study period, but use the results for only the first few years. This approach makes minimal assumptions beyond those of the analytical method itself. It also can be expensive. While EGEAS can be used this way, it also offers a less expensive method referred to as the extension period.

Most end effects (i.e., the causes of this distortion) can be put in one of two categories:

1. Capital left unused at the end of the planning period
2. Shifts in the relative economics of units caused by differences in escalation rates for various cost components

EGEAS uses an extension period to capture these effects in an approximate fashion. This period, which starts with the first year after the study period, may be finite or infinite in length, or may be omitted. EGEAS calculates capital and operating costs for the extension period, when included, and includes these costs in the total cost of an expansion plan. EGEAS uses present worth costs to avoid the infinite costs that would otherwise arise from inflation over an infinite extension period. The generation expansion optimization seeks to minimize the total present worth cost over both study and extension periods, so extension period modeling can greatly influence the optimization.

In modeling an extension period, EGEAS makes several basic assumptions:

System demand remains fixed at the level used for the last study period year. The rationale for this assumption is that EGEAS does not consider new resources in the extension period, and thus could not meet increased system demand.

The operating characteristics of all resources remain fixed at the values used for the last study period year. Included are such parameters as capacity, forced outage rate, maintenance requirements, and energy limits. This assumption itself can lead to distortions if resource characteristics vary widely from year to year. For example, EGEAS would tend not to install a generic unit with a high immature forced outage rate in the last study period year because the resource would retain the high forced outage rate throughout the extension period.

Operating costs continue to escalate in the extension period, and escalation rates can vary from year to year. Components of the operating cost include fuel costs, fixed and variable O&M costs, purchased energy costs, and unserved energy costs.

Non-EGEAS assets values (e.g., book depreciation, rate base value, CWIP balance, operating cost, and other electric revenues) continue to change in the extension period if a user supplies new values for those years. However, users only are allowed to supply values for a limited number of years and there is no escalation beyond that point, so eventually the non-EGEAS assets remain constant. The non-EGEAS assets values in the extension period are used only for calculating system average rate.

Capital fixed charges are applied each year until the resource reaches the end of its book life.

Any resource that retires during the extension period is immediately replaced by a resource with the same operating characteristics as those of the original unit in the last study period year. (For calculating operating costs, this assumption is the same as assuming that the original unit never retires.) If the capital installation costs are known, EGEAS calculates the capital fixed charges for the replacement resources, considering cost escalation between the installation years of the original and replacement resource. A replacement resource has the same operating life as the original unit it replaces. Thus, for a long extension period, the replacement resource may also retire. EGEAS keeps installing additional resources as necessary. In particular, an infinite extension period requires an infinite stream of replacement resources for each original resource. EGEAS replaces retiring resources whose capital costs are unknown, but makes no attempt to account for their capital costs.

If the capacity of a resource increases over time with corresponding new streams of capital fixed charges, both the capacity increases and the new capital expenses cease at the end of the study period. When such a unit is replaced in the extension period at the end of its operating life, the capital expenses for the replacement unit follow the same pattern as is used for the original unit. In particular, the capital expenses are calculated assuming the same capacity increases as for the original unit with no further increases after the point where the original unit hit the end of the study period. For production costing, however, EGEAS assumes that the capacity of that replacement unit is constant.

EGEAS calculates capital fixed charges and fixed O&M costs for each year of the extension period, discounts them to the base year, and sums them to obtain extension period totals. For a long period, the cost escalation rates eventually stabilize at a fixed rate (because you can only supply a limited number of different rates). Thus, for an infinite extension period, EGEAS calculates the costs using the formula for the sum of a geometric series.

Calculating variable operating costs (i.e., fuel, variable O&M, purchased energy, unserved energy) for a long extension period would be prohibitively time consuming if a separate production costing dispatch were performed for each year. Thus, EGEAS dispatches the system for just one year and projects the costs for the entire extension period. The subperiod detail for this one year is the same as in the last study period year. For each unit, EGEAS escalates each cost component, and constructs a time-weighted sum that represents the total present worth cost for a single MWH in every year of the extension period. For example, consider the following simple case:

Base year: 1988 **Extension period: 1998 to 2000** **Discount rate: 10 percent**		
Operating costs for unit A		
1998	50.00	$/MWH (current dollars)
1999	53.00	
2000	56.00	

For this example, EGEAS would calculate the time-weighted present worth cost for unit A as:

$$\frac{50.00}{(1.10)^{10}} + \frac{53.00}{(1.10)^{11}} + \frac{56.00}{(1.10)^{12}} = 55.70\,\$/\text{MWH}$$

where the denominators represent discounting by 10 percent back to the base year (1988). The final cost ($55.70) represents the total cost in 1988 dollars of producing 1 MWH in 1998, 1 MWH in 1999, and 1 MWH in 2000.

Input Requirements

The input requirements for analyzing expansion plans fall into four main categories: (i) system load; (ii) existing resources; (iii) new resources; and, (iv) modeling parameters.

System Load

For the system load, model users must supply peak loads, annual energies, and hourly load shapes for each year of the study period. Load shapes are input as load values at hourly or 5-minute intervals for the entire year. EGEAS does not contain a load forecasting model.

Existing Resources

For predetermined resources, users must supply the operating characteristics and costs for each generating facility that has already been built (existing resource), or is already planned (committed resource). The basic generating unit characteristics include capacity, forced outage rate, maintenance requirements, heat rate, energy limitation (hydro and storage units), and storage efficiency. Operating costs include fuel costs, and fixed and variable components of the O&M cost. For the basic optimization problem, capital costs are optional for committed and existing resources. However, when they are supplied, they have no effect on the optimization because they are sunk costs, and thus the same for all possible expansion plans. However, for certain economic feasibility constraints, users must supply capital costs for the committed resources and existing resources, or include them in the “non-EGEAS” rate base values supplied for each year.

New Resources

For new capacity options, users must supply basically the same information as for predetermined resource, including operating characteristics and costs. The operating characteristics for a new resource usually are less precise than those for an existing resource because the resource has no operating history, and may not be fully designed. Capital cost information is required for a new capacity option because the decision to include a new unit depends on its capital as well as its operating costs.

Modeling Parameters

For modeling parameters, users must supply options that define the problem to be solved, the methods to be used, and the level of modeling detail. These include economic assumptions, such as base year, price elasticity, and discount rate. For example, users must define study and extension periods, provide subperiod definitions, select an optimization method, choose possible planning alternatives, and set constraint bounds. Some modeling parameters are used only for specific optimization methods.

3
PRODUCTION COSTING

Overview

Variable generating costs can account for more than half of the total cost of an expansion plan. Therefore, accurate calculation of production costs is essential in determining the least cost expansion plan; however, long-term optimization requires that generating costs be calculated for many different expansion scenarios, so computational efficiency is also important. Thus, generation expansion and integrated resource planning requires a production costing method that is both accurate and efficient.

Chronological hourly production costing models can yield accurate production costs, but the computational effort is too great in the context of long-term optimization. On the other hand, probabilistic methods are much faster, but are not as accurate. To resolve this problem, EGEAS uses a probabilistic production costing method based on the Booth-Baleriaux algorithm.

The method used by EGEAS is capable of handling a wide variety of resources including thermal (nuclear, fossil steam, combustion turbine, combined cycle), hydro (reservoir hydroelectric, limited energy plant), storage (pumped hydro, battery, compressed air, superconducting magnets), and non-dispatchable technologies (solar, wind, cogeneration). The methodology features modeling of partial outage states, automatic maintenance scheduling, must-run resources, spinning reserve, fuel budgeting and a wide-array of cost inputs. Demand-side management resources are also featured in EGEAS. Costs of DSM technologies including costs, rebates, administrative and marketing cost, and T&D benefits can all be modeled. In addition, EGEAS allows modeling of purchase and sale contracts, interconnections to neighboring utilities or power pools, and environmental emissions and allowances. EGEAS uses basically the same production costing methodology for the Generalized Benders, dynamic programming, and pre-specified pathway analysis options.

Load Model

For production costing, EGEAS uses a load duration curve representation for system demand. A load duration curve (LDC) is the complement of the cumulative probability distribution for the hourly loads. Figure 3-1 shows an original LDC constructed from a set of hourly loads and an inverted LDC, where the load and probability axes are interchanged. EGEAS uses the inverted form for production costing. For more information about LDC's and how they are constructed, refer to the System Load section of Chapter 2.

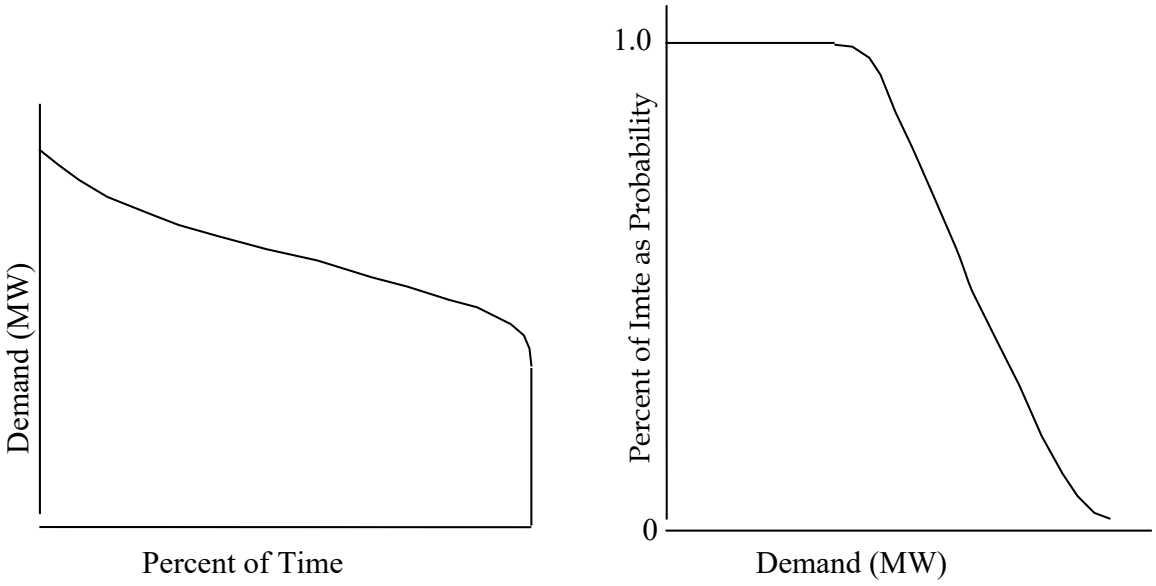

a. Original Load Duration Curve b. Inverted Load Duration Curve

Figure 3-1
Conversion of Load Duration Curve to Cumulative Probability Curve

While the exact shape of the LDC depends on the underlying distribution of loads, some general features are constant. First by way of definition, the height of the curve at any point gives the probability that the load exceeds that megawatt value. The LDC always has a value of 1.0 at load level zero and stays at that value up to the minimum load level. The point where the LDC begins to decrease is the minimum load point (sometimes called the “knee point”). All hourly loads must equal or exceed this load level. Beyond this point the LDC decreases monotonically until it reaches the value of 0.0 at the peak load. Beyond the peak load, the LDC stays at 0.0 because the load can never exceed its maximum value. When multiplied by the number of hours in the time period, the area under the curve gives the total amount of energy needed to be served (in MWH).

EGEAS gives you a choice of two methods for representing load duration curves: piecewise linear and analytic. A description of these two representations is given at the end of this chapter after the basic production costing algorithm has been presented.

Resources Model

The production costing model in EGEAS uses four basic types of resources to meet the system load: thermal, hydro, storage, and non-dispatchable technology (NDT). Each of these types can represent either supply-side resources (generating units) or demand-side resources (DSM programs). Because of differences in their characteristics and limitations, the four types of resources are modeled differently. The Probabilistic Production Costing Algorithm section below describes how thermal resources (both supply-side and demand-side) are dispatched. The Limited Energy Plants and Storage Units sections extend the algorithm for hydro and storage resources. NDT resources are modeled by modifying the LDC before the other types of resources are dispatched. The algorithm used for NDTs is described in the Non-Dispatchable Technologies section.

Key parameters used for production costing include the operating capacity (MW), the forced outage rate, maintenance requirements, heat rate, fuel type, and variable O&M cost. Thermal and hydro resources can have multiple loading blocks to represent partial outage states and heat rate curves. Refer to the Resources: Supply-Side and Demand-Side section of Chapter 2 for more information on how EGEAS represents resources.

In addition to the four basic types of resources, EGEAS can handle contributions from purchase and sale contracts and from interconnections with other utilities or power pools. The modeling of purchase and sale contracts is described in a later section of this chapter while the description of interconnections is in Chapter 8.

Probabilistic Production Costing Algorithm

This section describes the basic algorithm EGEAS uses for performing production costing with thermal resources. First, the algorithm is given for single-block resources and then it is extended to multiple-block resources. Later sections extend the algorithm to hydro and storage resources.

Single-Block Thermal Resources

For production costing, EGEAS uses a probabilistic method based on the Booth-Baleriaux algorithm. System demand is in the form of a load duration curve (LDC). System resources are treated as random variables with capacities and forced outages rates. In the simplest case discussed here, each resource has a single block of capacity with a single forced outage rate, which measures the fraction of time that the resource is out of service due to mechanical failure. The complement of the forced outage rate is availability, which gives the fraction of time that the resource can contribute towards meeting the load. As an example, consider a 100-MW unit with a forced outage rate of 0.2. Eighty percent of the time, this unit could supply 100 MW and twenty percent of the time, it could supply no power. A basic assumption is that the outages of one resource are independent of the outages of all other resources. The treatment of maintenance requirements is left to a later section.

Resources are loaded under the LDC according to a merit order that reflects increasing marginal operating costs ($/MWH). Resources near the beginning of the merit order operate for longer periods of time than do later units. Thus, production cost is minimized by maximizing the operation of the least expensive resources. The Dispatch Order section provides more detail on how resources are ordered.

To illustrate how resources are loaded, consider first the simple case where all resources are completely reliable (forced outage rates are all zero). In this case, resources are loaded as shown in Figure 3-2. Each shaded box represents the loading of a single resource. The width of the box is the capacity of the unit and the area, when multiplied by the number of hours in the time period, yields the MWH output of the resource. The average height of a box is the capacity factor for the resource, ranging between 0.0 and 1.0. As long as the total system capacity equals or exceeds the peak load, there is no unserved energy, the sum of the resource energies equals the energy for the initial load.

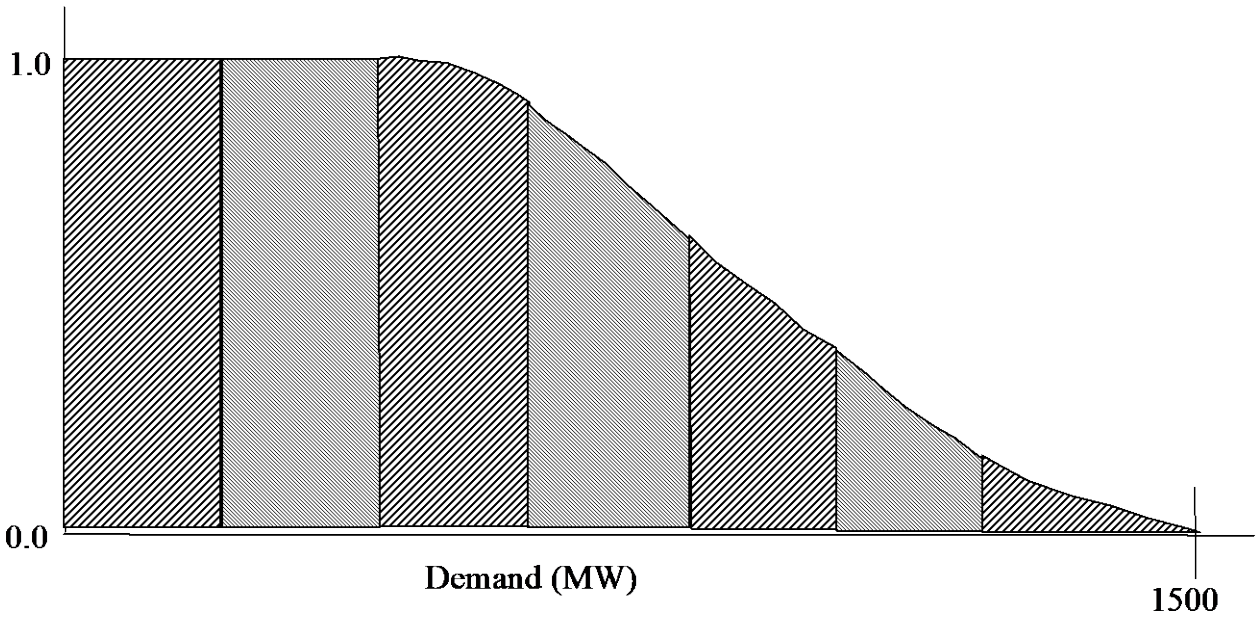

Figure 3-2
Loading Resources Without Outages

Generating units and DSM programs, however, are not completely reliable. When a resource is forced out due to some mechanical outage, resources later in the dispatch order are called on to increase their output to make up for the lost generation. Each succeeding unit is loaded to meet the "equivalent load", which is the sum of the initial system demand and the forced outages of the previously loaded units. After the resource is loaded, its forced outages are convolved onto the equivalent load duration curve (ELDC) to derive the new ELDC, which reflects the demand seen by the next resource in the merit order. The convolution process is illustrated in Figure 3-3. The first curve (a) shows the initial LDC for a period of 168 hours (one week); the area under the curve, when multiplied by 168, gives the total system energy to be met. The shaded area represents the energy generated by a 500-MW unit, assuming it operates 100 percent of the time. This energy is:

$$(500\,\text{MW}) * (168\,\text{HR}) = 84{,}000\,\text{MWH}$$

If the unit has a 20-percent forced outage rate, however, the shaded area must be multiplied by the availability (0.8) to obtain the expected energy generated:

$$(84{,}000\,\text{MWH}) * (0.8) = 67{,}200\,\text{MWH}$$

The load seen by the second unit depends on whether the first unit operates. If the unit operates, the full 84,000 MWH is generated and the load for the next unit is represented by the area under the original LDC beyond the 500-MW point; this situation is illustrated by the second curve (b) in Figure 3-3. If the first unit fails, no energy is generated and the load for the next unit is represented by the entire area under the initial LDC. Because the next unit will be loaded to the right of the 500-MW point, the initial curve is shifted to the right by 500 MW, the capacity of the first unit; this situation is illustrated by the third curve (c) of Figure 3-3. The final curve (d) in Figure 3-3 shows the curve that results from summing the second and third curves, weighted by the probabilities of the first unit being available (0.8) and being forced out (0.2), respectively.

This final curve is the ELDC under which the second unit will be loaded. The process of constructing the last curve (d) from the original curve (a) is known as convolution.

Mathematically, the convolution process is given by the formula:

$$G'(L) = p * G(L) + q * G(L - x)$$

where:
- G(L) = initial LDC at load level L
- x = capacity of unit
- q = forced outage rate of unit
- p = availability of unit (1 - q)
- G'(L) = equivalent LDC after convolving unit

Several important pieces of information can be obtained from the equivalent load duration curve; these are illustrated in Figure 3-4. First, the point at which a unit is loaded is known as the load (or loading) point. Because each unit is loaded immediately after its predecessor, the load point is the sum of the capacities of all units loaded before the current unit. The capacity point indicates the end point for loading of a unit. The difference between the load point and the capacity point for a unit is that unit's capacity. (A hydro or storage unit may be loaded at less than full capacity, in which case the capacity point does not reflect the full capacity of the unit.) The capacity point for the last loaded unit measures the total capacity of the system (point B in Figure 3-4).

The last point on the initial LDC, the point where the curve goes to zero, is the system peak load. When a generating unit is convolved, this last point is shifted to the right by the capacity of the unit. This shift represents the possibility that the unit fails to operate; the next unit is loaded at the capacity point of the first unit and there is a possibility that it sees the entire system demand. Thus, the distance from a unit's capacity point to the last point of the ELDC after it is convolved is always the initial system peak load. In Figure 3-4, the distance from point B to point C is the system peak load. After the last unit is loaded, the capacity point (B) is the total system capacity and the last point on the ELDC (C) is the sum of the system peak load and the system capacity.

After a unit is loaded, the height of the ELDC at the unit's capacity point gives the probability that the system demand exceeds the available capacity. In particular, after the last unit is loaded, this height gives the probability (or fraction of time) that the entire system load cannot be met. Thus, the height of the ELDC at the capacity point of the last unit gives the loss of load probability (LOLP). When multiplied by the number of hours in the time period, the area under the ELDC beyond the capacity point of a unit gives the amount of energy remaining to be served. After the last unit is loaded, this energy is the expected unserved energy. The expected unserved energy never goes to zero when the generating units have forced outages because there is always a possibility, however small, that all units fail simultaneously (because of their independence).

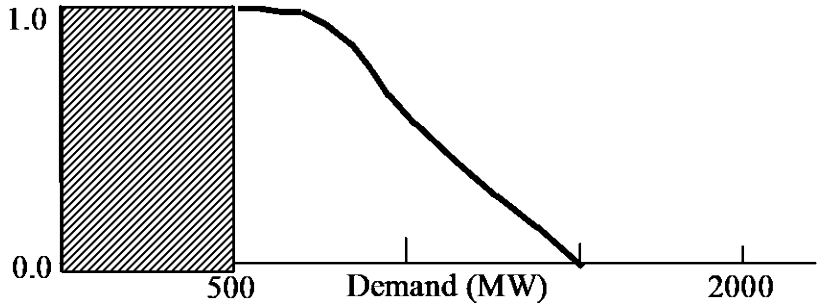

a. Loading of a 500 MW Unit

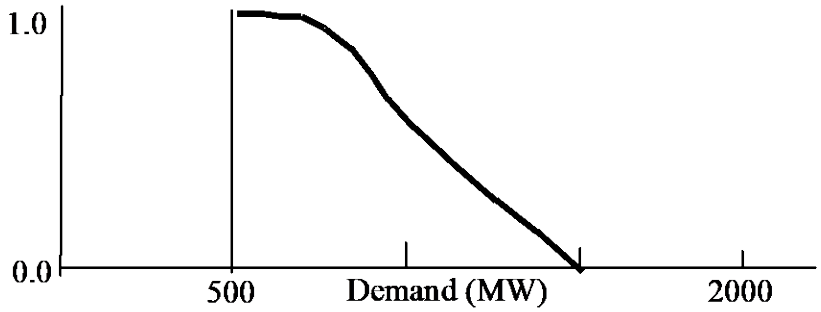

b. Load After First Unit Operates

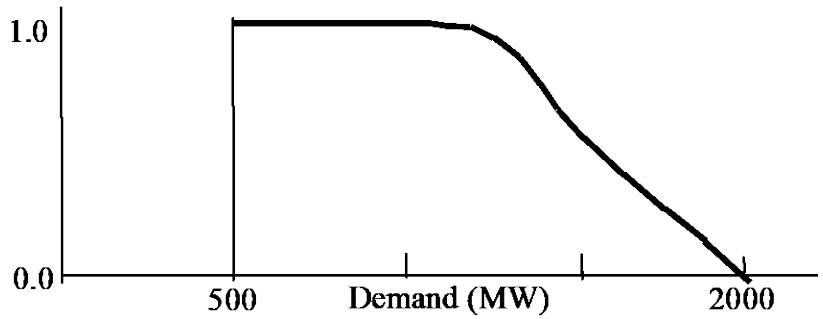

c. Load After First Unit Fails

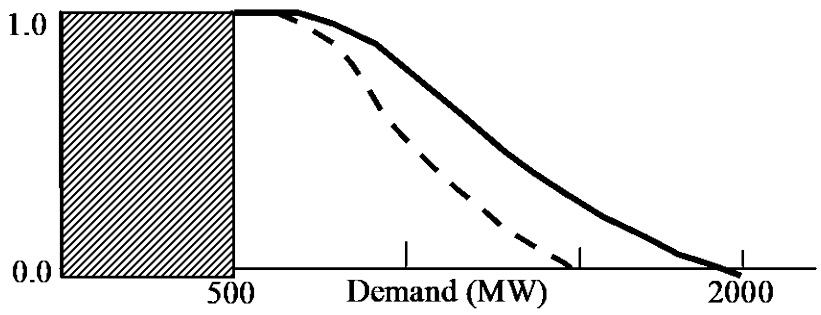

d. Equivalent Load Seen by Second Unit

Figure 3-3
Graphic Illustration of Convolution

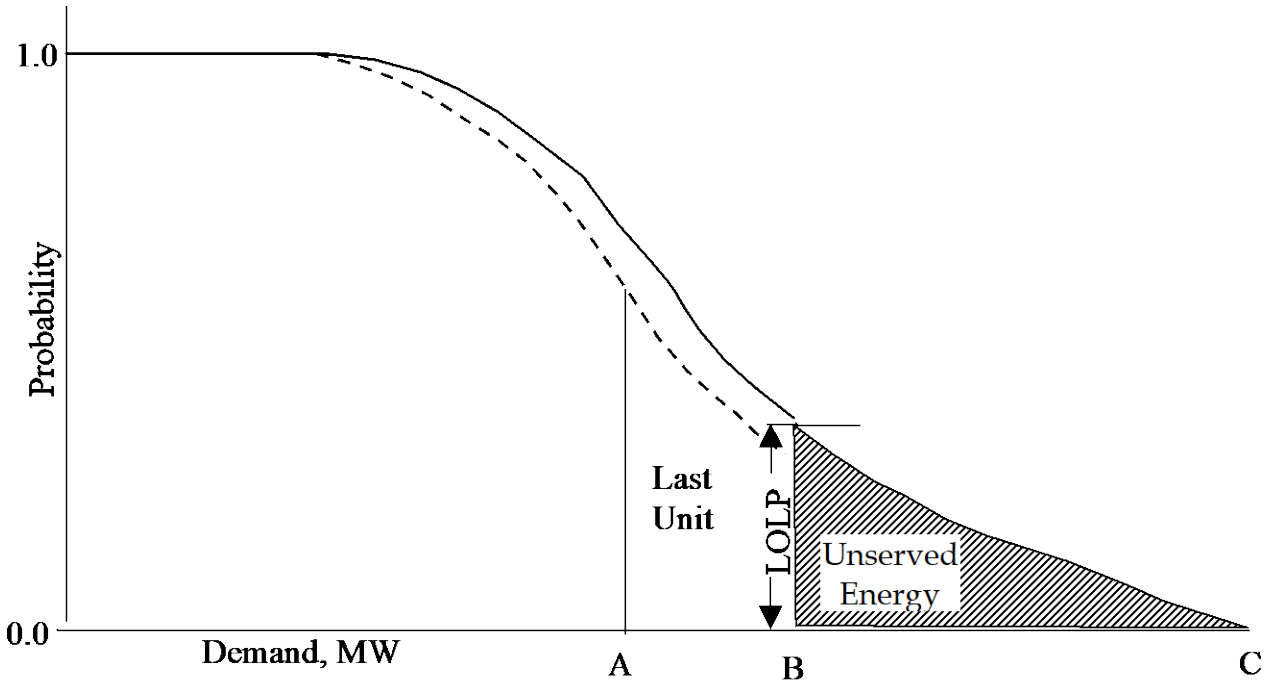

A = Load Point for Last Unit

B = Capacity Point for Last Unit

C = Maximum Equivalent Load

Figure 3-4
LOLP and Unserved Energy

The energy generated by a unit can be calculated in one of two ways. With the direct method, described earlier in the discussion of convolution, the area under the ELDC between a unit's load and capacity points is multiplied by the unit's availability to obtain the expected energy generated; the area under the tail of the ELDC after the last unit yields the expected unserved energy. Because of numerical imprecision in the area calculations, the sum of the generation and expected unserved energy may not exactly match the initial system energy. With the energy balance method, the energy generated by a unit is the difference between the unserved energy remaining before and after the unit is loaded. Note that the unserved energy before a unit is loaded is the same as the unserved energy after the previous unit was loaded. Thus, only one new area calculation is needed for each generating unit. EGEAS uses the energy balance method because it ensures that the sum of the energies generated and the remaining unserved energy equal the initial system energy.

Multiple-Block Thermal Resources

When a thermal resource is modeled with multiple loading blocks, EGEAS loads each block separately. These blocks may be loaded in contiguous positions or may be separated by blocks of other resources, depending on the relative marginal costs. Because the outage of one block prevents all higher blocks from operating, EGEAS cannot treat the blocks of a resource independently. Instead, EGEAS must consider the joint probability distribution. For a resource with n blocks, convolution of the outages onto the LDC is given by the formula:

$$G'(L) = p * G(L) + \sum_{i=1}^{n} q_i * G(L - X_i)$$

where:

$G(L)$ = initial LDC at load level L

q_i = partial outage rate for block i

x_i = capacity for blocks i through n

p = availability for the entire resource (1 - sum of q_i)

$G'(L)$ = equivalent LDC after convolving resource

When the first block of a multiple-block resource is loaded, EGEAS computes the energy generated by that block and convolves its outages onto the ELDC using the formula for a single block. When the second block is loaded, EGEAS must account for the dependence of block 2 on block 1. Three steps are required for this process. First, EGEAS removes the outages of the first block from the ELDC; block 2 cannot operate when block 1 is forced out, so block 2 cannot generate to meet the block 1 outages. Second, EGEAS calculates the energy generated by the second block. Third, EGEAS convolves the outages of the first two blocks onto the ELDC using the formula for two blocks. Note that this formula for two blocks does not yield the same results as convolving the two blocks separately using the formula for a single block. In a similar manner, EGEAS loads a higher block n (n greater than 2) by first removing the outages of the first n-1 blocks, then computing the energy generated by block n, and finally convolving the outages of the first n blocks.

Removing a block's outages is accomplished through deconvolution, which simply reverses the convolution process. For a resource with n blocks, deconvolution of the outages from the LDC is given by the formula:

$$G(L) = \frac{G'(L) - \sum_{i=1}^{n} q_i * G(L - x_i)}{p}$$

where G' is the ELDC before deconvolution and G is the ELDC after deconvolution. To illustrate the process of loading a multiple-block resource, consider the following example:

Order of blocks to be loaded:

Resource A, block 1 (block A1)

Resource B

Resource A, block 2 (block A2)

Resource C

Resource A, block 3 (block A3)

Procedure:

1. Start with initial LDC #0.
2. Compute the energy generated by block A1 (use area under LDC #0 for the load point to the capacity point of A1).
3. Convolve the outages of A1 to obtain LDC #1.

4. Compute the energy generated by resource B (use area under LDC #1).
5. Convolve the outages of B to obtain LDC #2.
6. Deconvolve the outages of block A1 to obtain LDC #3.
7. Compute the energy generated by block A2 (use area under LDC #3).
8. Convolve the outages for blocks A1 and A2 onto the LDC, using the formula for 2 blocks,
   to obtain LDC #4.
9. Compute the energy generated by resource C (use area under LDC #4).
10. Convolve outages of C to obtain LDC #5.
11. Deconvolve the outages of blocks A1 and A2, using the deconvolution formula for 2 blocks, to obtain LDC #6.
12. Compute the energy generated by block A3 (use area under LDC #6).
13. Convolve the outages for blocks A1, A2, and A3 onto the LDC, using the formula for 3 blocks, to obtain LDC #7.

In addition to the initial load, the various LDCs contain the outages of the following resources and blocks:

LDC	Outages of
0	none
1	A1
2	A1, B
3	B
4	A1+A2, B
5	A1+A2, B, C
6	B, C
7	A1+A2+A3, B, C

Dispatch Order

EGEAS minimizes the production cost by using a merit order that reflects increasing operating costs. These basic operating costs are fuel cost, variable operating and maintenance (O&M) cost. Detailed costs (e.g., DSM customer costs, transmission/distribution costs) are included when they are modeled as variable O&M costs. All costs are expressed as dollars per megawatt hour ($/MWH). When a resource uses two fuels, EGEAS uses a weighted average of the two fuel costs, based on the target percentage for use of the secondary fuel. (For further details, see the section on Dual Fuels and Fuel Constraints). EGEAS can also include adjusters or costs for emissions to alter the dispatch order, see Chapter 9 for more information.

EGEAS provides you with considerable flexibility in determining which costs are included when ordering resources. There are five options:

Weighted average fuel cost plus variable O&M (default)

Weighted average fuel cost only

Variable O&M cost only

Primary fuel cost plus variable O&M

Primary fuel cost only

Detailed costs modeled as variable O&M are included whenever the basic variable O&M costs are included.

The dispatch order is also affected by such modeling features as dispatch modifiers, must-run units, and spinning reserve. The sections below discuss these effects as well as the treatment of multiple-block resources.

Loading Blocks

When loading blocks are modeled, EGEAS divides generating units (or DSM resources) into several blocks of capacity that can be dispatched separately. Each block can have a different heat rate, which results in a different fuel cost for each block. Variable O&M costs are assumed to be the same for all blocks of a given unit. Before EGEAS orders the individual loading blocks, some adjustments in the marginal operating costs must be made. Because EGEAS uses the average heat rate for the first block, the no load costs for the unit are included in this first block. Higher blocks use average incremental heat rates, which do not include any no load costs. Thus, the marginal operating cost for the first block of a unit is often higher than the costs for higher blocks. If EGEAS ordered blocks strictly on the marginal operating cost, the higher blocks of a unit could be loaded before that unit's first block.

To keep blocks of a given unit in the proper order, EGEAS adjusts the marginal cost used for determining the dispatch order. Note, however, that the true costs are still used when calculating the production cost. The adjustment procedure looks at the weighted average cost of two or more adjacent loading blocks. If this average cost is less than the cost of the first block in the group, then the average cost replaces the true cost for all blocks within the group. The result is that these blocks get loaded together. In essence, this procedure causes units to be loaded in the same manner as when loading blocks are not modeled; however, using separate costs for the individual blocks does result in a more accurate production cost. In addition, loading block modeling makes possible the modeling of must-run units and spinning reserve.

Must-Run Units

Often utilities must operate certain units all the time, even when they are not the least costly units available. Some units are not designed for cycling and cannot be shut down and then restarted in a timely fashion. Other units are critical to system security (perhaps because they are physically close to the customer load) or are needed to maintain voltage control. To model such units (or DSM resources), EGEAS has a must-run unit capability. For modeling flexibility, hydro units can also be designated as must-run units. Essentially, EGEAS places must-run units first in the dispatch order, regardless of their marginal operating cost. If loading blocks are modeled, only the first block of a must-run unit is so ordered; the higher blocks are loaded according to economics. Thus, the individual blocks of a must-run unit are not loaded contiguously. When

there are several must-run units or blocks of units, EGEAS orders them according to their true marginal operating costs.

The must-run feature can also be used to alter the way multiple-block units are dispatched. As described in the previous section on Loading Blocks, all the blocks of a multiple-block unit are often loaded contiguously with an average cost, because of the no-load costs associated with the first block. If a unit is designated as must-run, however, the first block is placed at the beginning of the dispatch order and the remaining blocks can be loaded according to economics without the adjustments for the first block. To illustrate this process, consider two identical units (A and B) with the following characteristics:

Block	Capacity	Marginal Cost	Adjusted Marginal Cost
1	100 MW	50 $/MWH	38 $/MWH
2	100	30	38
3	100	34	38
4	100	60	60

If these units are not designated as must-run, the order in which the blocks will be loaded is as follows:

A1, A2, A3, B1, B2, B3, A4, B4

For ordering purposes, the first three blocks of each unit have a cost of 38 while the fourth block has a cost of 60. When there are units with identical costs, all the blocks of one unit are placed before the blocks of the other unit(s).

If these units are designated as must-run, the order in which the blocks will be loaded is as follows:

A1, B1, A2, B2, A3, B3, A4, B4

The higher blocks are loaded according to economics with no need for adjustment because the first blocks have been separated because of the must-run status. Thus, both second blocks (with a cost of 30) are loaded before third blocks (with a cost of 34). If you want to achieve this type of loading, but do not want the first block operating all the time, you can make the first block very small, but give it all the no-load costs. Then only a small portion of the unit is operating with a high capacity factor.

Because EGEAS places must-run units first in the dispatch order, they are usually loaded under the portion of the load duration curve where the height is 1.0 (below the minimum load point). Such units operate all the time, subject to their availability. If there are many must-run units, however, their capacity may exceed the minimum load so some of the units are not needed all the time. EGEAS gives you two options for handling such situations. If you request the dump energy option, then all must-run units are operated to their full capability; any energy that cannot be used to meet the system load is dumped. The cost of producing the dump energy is included in the production cost. If you do not request the dump energy option, then must-run units are operated only when needed to meet the system load.

One additional option is a dump penalty factor, which is used to determine an additional cost to be applied when there is dump energy. The penalty factor is multiplied by the full cost of generating the energy that is being dumped. This cost is then added to the production cost. For

example, if the cost of generating the dump energy is $10 million and the penalty factor is 1.5, and additional $15 million will be added to the production cost, meaning that $25 million total was spent for the energy that was dumped.

Spinning Reserve

Utilities must generally provide backup capacity in one or more categories to handle emergency outage situations. Some of this capacity may be on-line running at less than full capacity, while other capacity may be off-line but ready to be brought on-line quickly. EGEAS is not designed to model all types of reserve. Instead, EGEAS concentrates on spinning reserve, additional capacity from units operating at less than full capacity. Spinning reserve can be important for determining accurate production costs because it affects the dispatch order, which in turn affects the generation output of individual units. Only thermal resources can be used for spinning reserve.

To maintain spinning reserve, system dispatchers keep several units operating at less than full capacity. Thus, as the system demand increases, the dispatchers often bring several units on-line at their minimum capacity points before raising their output up to full capacity. Similarly, EGEAS loads the first blocks of several units before loading the corresponding higher blocks; the higher blocks provide the spinning reserve. Note that loading block modeling is a prerequisite for spinning reserve modeling.

When spinning reserve is modeled, you must specify the amount of reserve required. The amount may be specified in one of four ways: as a percentage of the system peak load, as a percentage of the capacity of the largest unit on the system, as a percentage of the total capacity of the two largest units, or as a fixed megawatt value. For the first three options, the values may exceed 100 percent. The spinning reserve requirement may be changed from year to year, or change from segment to segment.

Some technologies are more suitable for spinning reserve than others because of their cycling capabilities and high ramp rates. Thus, you must specify which units are available to provide spinning reserve. You can change a unit's spinning reserve status from year to year. In addition, you can impose a limit on the maximum spinning reserve contribution from any unit. This limit is specified as a percentage of the total capacity of the unit. For example, if a 600 MW unit has a first block capacity of 200 MW and the maximum contribution is 100 percent, then that unit can contribute 400 MW. However, if the maximum contribution is 60 percent, then the unit can contribute only 360 MW (60 percent of 600 MW). Note that the percentage applies to the total capacity of the unit, not just to the capacity of the higher blocks providing the reserve. The same percentage applies to all units and to all periods.

When setting up the dispatch order, EGEAS starts by ordering the loading blocks according to their adjusted marginal operating costs. (The adjustments to keep blocks of a particular unit in order are described in the **Loading Blocks** subsection). EGEAS places the first blocks of any must-run units at the beginning of the dispatch order. Next, EGEAS computes how much spinning reserve is provided by the higher blocks of whatever must-run units the user selected for spinning reserve. If the reserve requirement is not satisfied by just the must-run units, then EGEAS proceeds through the initial dispatch order, searching for the first blocks of units available to provide spinning reserve. As appropriate blocks are found, EGEAS places them, in order, after the must-run units. When the required spinning reserve is achieved, the process stops.

The process described above provides spinning reserve at low load levels. As the load increases, however, the higher blocks providing the reserve are loaded and eventually the spinning reserve goes to zero. EGEAS tries to maintain the reserve for as long as possible by delaying the loading of the higher blocks or by loading the first blocks of additional units. EGEAS provides you with three options for this part of the algorithm.

For the first option, EGEAS dispatches the remaining loading blocks based only on economics. If much of the spinning reserve is provided by base or intermediate units, then the reserve will be small or zero at higher load levels. If a significant portion of the reserve is provided by peaking units, however, then the reserve will be retained through most load levels.

For the second option, EGEAS places the higher blocks of the spinning reserve units last in the loading order. This option essentially maintains the spinning reserve until the last possible moment. If a significant portion of the reserve is provided by base units, this option may not be realistic because inexpensive units are under-utilized.

For the third option, EGEAS changes the units providing spinning reserve as the load level increases. Essentially, when EGEAS loads a block that is providing reserve, it checks whether the remaining reserve will be adequate. If not, EGEAS searches higher in the dispatch order for an additional spinning reserve unit; if it finds one, it moves the first block before the current block so this new unit's higher blocks provide reserve to replace what is being lost. Thus, less expensive units provide reserve at low load levels and more expensive units provide reserve at high load levels. In most cases, this third option comes closest to matching actual system dispatching.

The following simple example illustrates how spinning reserve is modeled with the three options:

Spinning reserve requirement: 120 MW		
Units that can provide spinning reserve:		
	Block	**Capacity**
SA	1	100 MW
	2	50
	3	70
	4	80
SB	1	50
	2	50
SC	1	50
	2	100
Units that cannot provide spinning reserve:		
A,B,C,D		
Economic dispatch order:		
A		
SA#1		
SA#2		
SA#3		
B		
SB#1		

SB#2
C
SC#1
SC#2
D

For option 1, the first block of unit SA is moved to the beginning of the dispatch order and its higher blocks (2 through 4) provide 200 MW of spinning reserve, sufficient to meet the requirement of 120 MW. From this point on, the units are loaded according to the initial (economic) dispatch order. Unfortunately, the spinning reserve is lost almost immediately as the higher blocks of SA are loaded after unit A, as shown below:

Option 1	**Available**
Dispatch Order	**Spinning Reserve**
SA#1	200 MW
A	200
SA#2	150
SA#3	80
SA#4	0
B	0
SB#1	0
SB#2	0
C	0
SC#1	0
SC#2	0
D	0

For Option 2, the first block of unit SA is also moved to the beginning of the dispatch order, but this time the higher blocks of SA are moved to the end of the dispatch order. Thus, the 200 MW of spinning reserve is preserved for all load levels until those three blocks of SA must be dispatched to meet the load. Unfortunately, this means that the upper blocks of SA, which have fairly low cost, are under-utilized.

Option 2	**Available**
Dispatch Order	**Spinning Reserve**
SA#1	200 MW
A	200
B	200
SB#1	200
SB#2	200
C	200
SC#1	200
SC#2	200

D	200
SA#2	150
SA#3	80
SA#4	0

For Option 3, the first block of Unit SA is again moved to the beginning of the dispatch order and its higher blocks provide 200 MW of spinning reserve. When block 2 of SA is dispatched, its spin (50 MW) is lost so the available reserve drops to 150 MW. However, this is still above the 120 MW requirement. Dispatching block 3 of SA would drop the reserve by another 70 MW, to 80 MW, which is below the requirement, so EGEAS moves the first block of SB before SA block 3. This move adds the 50 MW from block 2 of SB to the available spinning reserve so SA block 3 can then be loaded. At this point the available spinning reserve is 130 MW (80 MW from SA block 4 and 50 MW from SB block 2). The same process is repeated when EGEAS gets to SA block 4: SC block 1 is inserted before SA block 4 so the reserve from SC block 2 (100 MW) can make up for the lost reserve from SA block 4 (80 MW). When EGEAS gets to SB block 2, there are no more spinning reserve units, so the spinning reserve drops below the requirement.

Option 3	Available
Dispatch Order	**Spinning Reserve**
SA#1	200 MW
A	200
SA#2	150
SB#1	200
SA#3	130
SC#1	230
SA#4	150
B	150
SB#2	100
C	100
SC#2	0
D	0

Dispatch Modifiers

Many factors must be considered when dispatching generating units for an electric utility, including unit cycling capabilities, tie-line limitations, transmission losses, voltage control, fuel contracts, limited fuel supplies, and system security. Sometimes these factors dictate that units be loaded out of economic order. To help users model such situations, EGEAS allows the specification of a dispatch modifier for each unit. This modifier is added to the fuel and variable O&M costs for the unit to obtain the marginal operating cost used for setting up the dispatch order. For multiple-block units, the same modifier is added to the cost for each block. The modifier can be either positive (to delay the loading of a unit) or negative (to advance the

loading). The modifier can be changed from year to year. Note that the modifier is used only for ordering the units; it is not used in calculating the production cost.

Limited Energy Plants

This section deals with extending the basic production costing algorithm to limited energy plants. Included are supply-side hydro units and demand-side energy-limited resources. The basic feature of limited energy plants is that the output in gigawatt hours is limited in each time period.

Model Description

A limited energy plant (LEP) is similar to a thermal unit because it can be dispatched by the utility to meet system demand. As with a thermal unit, the output at any given time is limited by the unit's capacity and the unit is subject to forced outages and preventive maintenance requirements. An LEP is unlike a thermal unit, however, because its total generation over a period of time is limited by factors external to the plant. These factors include water availability (for driving a hydro or cooling a thermal unit), fuel availability (limited by contract or a specific market) and effects on the environment (release of pollutants or rise in the ambient temperature of a body of water). In EGEAS, LEPs are sometimes referred to as "hydro" units, whether or not the underlying technology is hydroelectric. For DSM resources similar concepts apply.

EGEAS assumes that LEPs, like all other units, are independent. In particular, fossil LEPs do not compete for limited fuel and hydroelectric plants on the same river system do not affect each other. Thus, you must specify a separate energy limit for each LEP. An LEP can be divided into loading blocks for more detailed modeling of outage states or heat rates. In such a case, the entire energy limit is assigned initially to the first block. If the first block does not generate up to the limit, the remaining energy is transferred to the second block. Similarly, the remaining energy after block 2 is transferred to block 3, and so forth until all blocks have been loaded or the energy limit is exhausted. The description of the methodology given below focuses on a single-block LEP, but applies with few changes to individual blocks of a multiple-block unit as well.

To satisfy an energy limitation, an LEP's output can be curtailed either by reducing the operating capacity or by reducing the unit's hours of operation. Whenever possible, EGEAS operates an LEP at full capacity to achieve the maximum effect on reliability (expected unserved energy); the energy constraint is met by limiting the hours of operation.

For production costing, an LEP is dispatched in much the same way as a thermal unit. In particular, the energy generated is the product of the area under the equivalent load duration curve (ELDC) and the unit's availability, and the outages for an LEP are convolved onto the ELDC. Before the unit dispatch begins, LEPs are placed in the merit loading order, according to their marginal operating costs, with the thermal units.

Methodology

During production costing, EGEAS proceeds through the merit loading order, dispatching each resource under the ELDC. In the case of an LEP, EGEAS loads the resource on a trial basis. If the energy generated is within the specified energy limit, then the LEP remains loaded; in essence, the LEP is treated just like a thermal resource. This situation, illustrated in the top curve (a) of Figure 3-5, can occur for any of the following reasons:

The energy limit is unrealistically high – greater than the product of the resource's capacity, availability, and the number of hours in the period.

The resource cannot generate up to the energy limit because it is on maintenance during all or a large portion of the period.

The resource is more expensive to operate than other resources on the system; when it is finally loaded, according to economics, only a portion of its energy is needed.

On the other hand, if the energy generated by an LEP during trial loading exceeds the resource's energy limit, then the LEP cannot remain loaded at that point. This case is illustrated in the bottom curve (b) of Figure 3-5. Such an LEP must be moved further out on the ELDC to a point where the generation does not exceed the energy limit. Thus, the loading of the LEP is "delayed" (loaded after other units with higher marginal operating costs).

EGEAS tries to dispatch each delayed LEP, at full capacity, where its output exactly matches its energy limit. Because the ELDC changes as the outages of each generating resource are convolved, EGEAS cannot determine in advance where to insert the delayed LEP in the dispatch order. Instead, EGEAS must check for the proper loading position after each succeeding thermal resource is loaded. Thus, the delayed LEP is loaded, on a trial basis, after each thermal resource. When the generation is within the energy limit, as illustrated in the top curve (a) of Figure 3-6, then the proper loading point has been found.

In general, the exact loading point will usually be in the middle of a thermal resource rather than at the boundary between two resources; in this case, the delayed LEP is said to "split" the thermal resource. The bottom curve (b) of Figure 3-6 illustrates a delayed LEP splitting a thermal resource. An invariance property states that the total energy generated by two adjacent resources is independent of the order in which they are loaded. Thus, EGEAS does not need to determine the exact loading point of the LEP within the thermal resource it splits. Instead, EGEAS determines the total generation by the two resources from the change in unserved energy. The generation by the LEP is set to the value of its energy limit and the generation by the thermal resource is the total generation less the LEP generation.

When several LEPs are delayed, they can compete for the same loading point. In such a case, the delayed LEPs may split the same thermal resource, and indeed split each other. Because of the energy invariance property, EGEAS sets the generation by each LEP to its energy limit, and the balance of the energy is allocated to the thermal resource. Note that when EGEAS reports the order in which resources were dispatched, delayed LEPs are always placed immediately before the thermal resource they split. This is simply a reporting convention.

For a delayed LEP with a small energy limit, the proper loading point can lie beyond the capacity point for the last loaded thermal resource. This case is illustrated in the top curve (a) of Figure 3-7. To prevent the resource from exceeding its energy limit, EGEAS loads the resource with a reduced capacity, as shown in the bottom curve (b) of Figure 3-7.

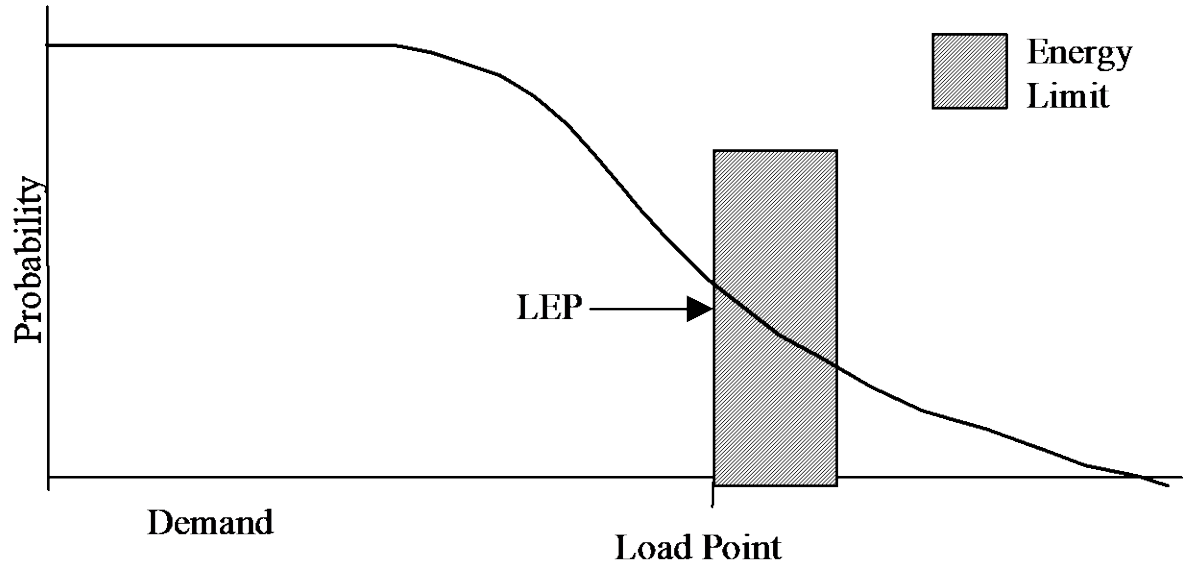

(a) LEP Can Be Loaded

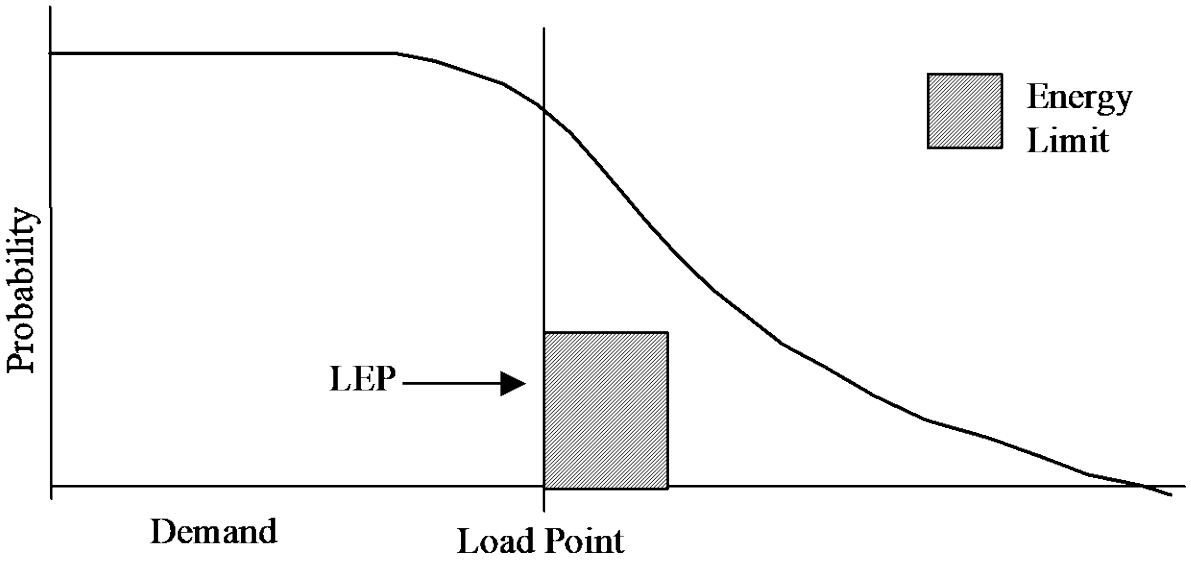

(b) LEP Must Be Delayed

Figure 3-5
Trial Loading of an LEP

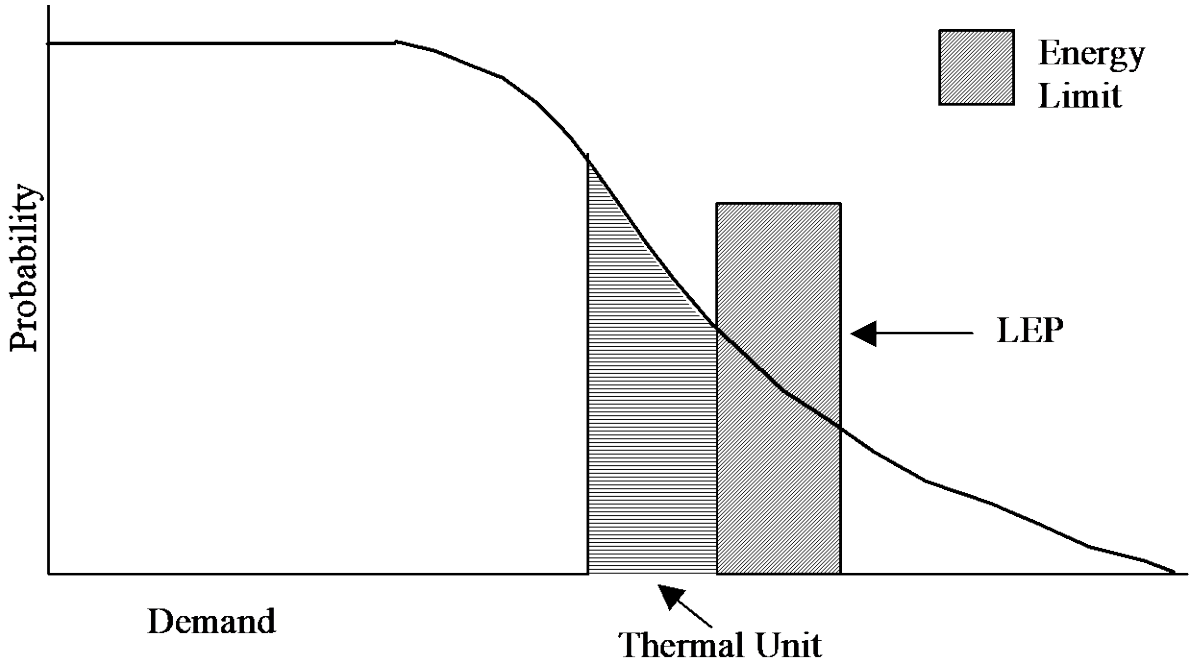

(a) Trial Loading After Thermal Unit

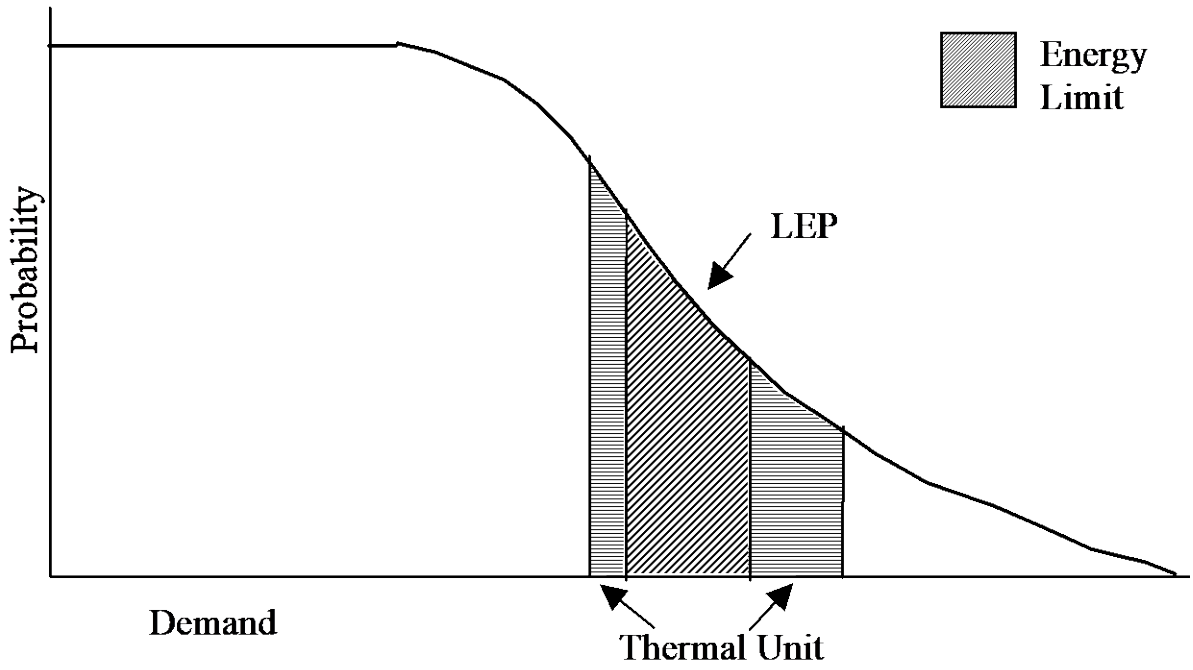

(b) Splitting a Thermal Unit

Figure 3-6
Loading of a Delayed LEP

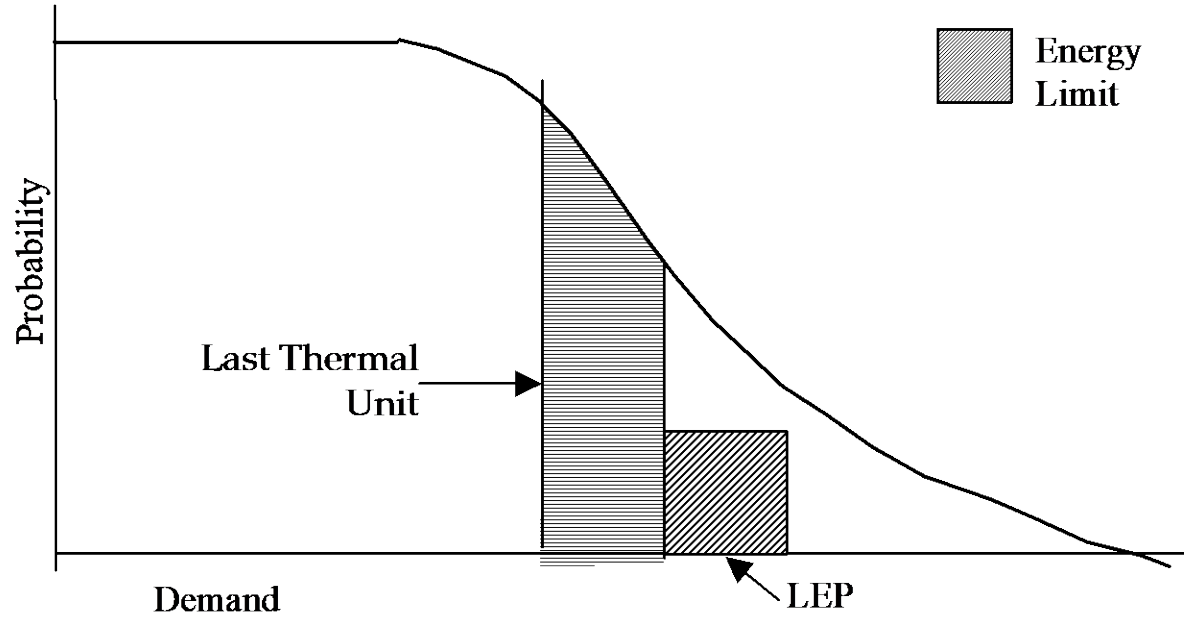

(a) Trial Loading of LEP At Full Capacity

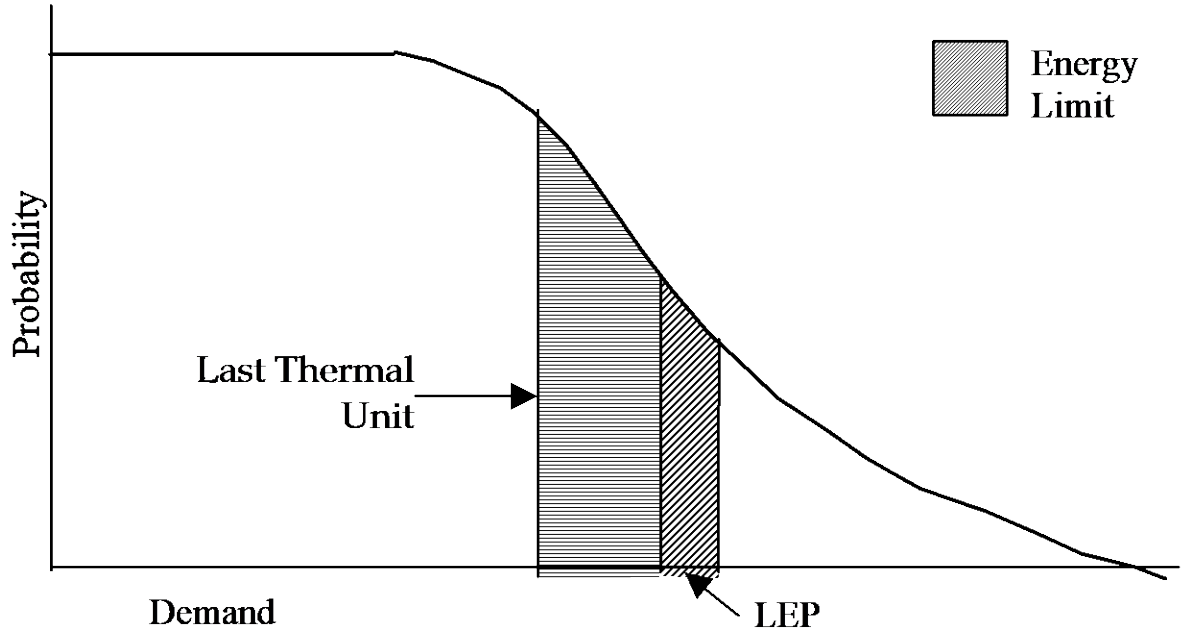

(b) LEP Loaded At Reduced Capacity

Figure 3-7
Delayed LEP Loaded at Reduced Capacity

Storage Units

The distinguishing characteristic of a storage unit is that it both takes energy from and supplies energy to the electric grid. The rationale for operating storage units is to store generation by inexpensive base-loaded units during off-peak load periods and return the energy during peak load periods to replace generation by expensive peaking units. To be economical, the difference in operating cost between the base units and the peaking units must be large enough to outweigh the energy losses and the costs of the storage operation.

Model Description

EGEAS can model several storage technologies including pumped hydroelectric, large batteries, compressed air, and superconducting magnets. For DSM, it can model commercial cool storage, heat storage, and other load shifting programs. However, there are some limitations:

Because EGEAS does not use a chronological model, some of the time-dependent limitations are not modeled precisely. The storage operation determined by EGEAS might in reality require charging when the reservoir (or other storage medium) was full or generating when the reservoir was empty. EGEAS does provide options for minimizing the chance of such situations.

EGEAS cannot model hybrid units – that can switch between hydroelectric and pumped storage operations.

EGEAS cannot model units that provide dedicated storage for a single unit, such as a solar photovoltaic installation.

EGEAS cannot account for losses that are dependent on the amount of time that the energy is stored (as with batteries).

EGEAS does not consider the dynamic benefits of operating storage resources, such as dependable peaking capacity, frequency regulation, load-following capability, and spinning reserve.

The basic operating characteristics of a storage resource are:

Generating capacity: the maximum instantaneous output of the storage resource.

Charging capacity: the maximum capacity for charging operations; can be different from the generating capacity.

Forced outage rate: the probability that the storage resource is unavailable because of mechanical failure. The same forced outage rate is used for both charging and generating operations.

Storage efficiency: the ratio of the energy generated to the energy charged. This efficiency incorporates losses from both the charging and generation operations. For most storage resources, the efficiency is less than 1.0; however, for compressed air units, which receive additional energy from burning fuel, the efficiency is greater than 1.0.

Energy limit: the maximum amount of energy that can be generated in a particular time period. The energy limit must account for the storage capacity of the resource (how much energy can be stored at any given time) and the number of operating cycles in the time period (e.g., is the resource charged and discharged once per day or once per week).

Charging energy limit: the maximum amount of energy that can be charged in a particular time period. EGEAS calculates the charging energy limit by dividing the (generation) energy limit by the storage efficiency.

Operating costs: include fixed and variable operating and maintenance (O&M) costs and fuel costs (see Chapter 2 for more details).

The user specifies all the above characteristics except the charging energy limit, which is calculated by EGEAS.

For improved computational efficiency, EGEAS allows you to group storage units (or resources) into storage projects. EGEAS then models storage operation at the project level. While serious inaccuracies may result from grouping dissimilar units, the results are more than adequate for long-term generation planning when the units grouped together have similar operating characteristics. Thus, the grouping of pumped storage units with similar efficiencies and energy limits (per unit of capacity) is recommended, but the grouping of units with different technologies (e.g., pumped hydro and compressed air) is not.

For a storage project, the charging capacity, generating capacity, energy limit, and charging energy limit are all sums of the corresponding quantities for the component storage units. The efficiency of the project is the ratio of the (generating) energy limit to the charging energy limit.

After production costing is complete, EGEAS allocates the storage energies to the component storage units in the same proportions as the unit energy limits. Consider two storage units (A and B) aggregated to form a project:

	Unit A	Unit B	Project
Generating capacity, MW	125	110	235
Charging capacity, MW	100	80	180
Generating limit, GWH	196	144	340
Charging limit, GWH	280	200	480
Efficiency	0.70	0.72	0.7083

In this example, the variables for Units A and B are supplied to the model, except the charging limit, which EGEAS calculates from the generating limit and the efficiency (e.g., 196/0.70 = 280 for Unit A). The project variables are sums of the corresponding variables for Units A and B, except the efficiency, which EGEAS calculates as the ratio of the generating limit to the charging limit (340/480 = .7083). If EGEAS determines that the storage project should generate 272 GWH, then this energy is distributed to Units A and B as follows:

Unit A: (196/340) * 272 = 156.8 GWH

Unit B: (144/340) * 272 = 115.2 GWH

Using the storage efficiencies, EGEAS determines the charging energies to be:

Unit A: 156.8/0.70 = 224.0 GWH

Unit B: 115.2/0.72 = 160.0 GWH

The total charging energy (224.0 + 160.0 = 384.0) matches the project total (272/.7083 = 384.0).

EGEAS divides the analysis of storage units into medium-term and short-term decisions. The medium-term decision determines how much energy should be generated (to maximize operational savings), while the short term decision determines how to dispatch the storage unit and how to provide charging energy. The two types of decisions are intertwined because the short-term dispatch algorithm affects the economics upon which the medium-term optimization is based, but the medium-term optimization determines the amount of energy for the short-term dispatch. For the medium-term decision, the user may instruct EGEAS not to optimize the generation amount; in this case, EGEAS sets the generation equal to the unit's specified energy limit. Bypassing the mid-term optimization this way will reduce computation time, and may be appropriate when the user can accurately predict the unit's future operation.

In this discussion of storage, the term "thermal" includes any LEP dispatched in economic loading order (i.e., not delayed) as well as all thermal units. Delayed LEPs are not included because their generation amounts are fixed at their energy limits.

Short-Term Decision

The charging and generating operations can be considered independently because the generation amount is already known and the amount of charging is readily determined by dividing the generation by the storage efficiency.

The generating side of a storage unit has the same modeling requirements as a single-block limited energy plant (LEP); the basic parameters are capacity, forced outage rate, energy limit, and operating cost. Note that the operating cost includes the costs associated with both the charging and generating operations; it does not include the costs for the thermal units providing the charging energy. Thus, EGEAS treats the generating side as an LEP. A description of the algorithm for dispatching LEPs is given in the **Limited Energy Plants** section.

The charging side of a storage unit is similar to the generating side in that its dispatch is governed by a capacity, forced outage rate, and energy limit. However, instead of satisfying any system demand, the charging side increases the equivalent demand which must be met by the remaining units on the system. Thus, the charging side is modeled as a "fictitious" LEP. The forced outage rate of this LEP is the availability of the storage unit (rather than its forced outage rate) because the charging side imposes an added load on the system only when it is available.

Before any generating units are loaded, the fictitious LEP is convolved onto the initial load duration curve (LDC), as illustrated in the top portion (a) of Figure 3-8. The energy added to the LDC is the energy required if the storage unit is charged all the time, subject to its availability. The charging side cannot use all this energy, however, because the unit cannot be in a charging mode all the time (and because it probably exceeds the charging energy limit). Thus, the unneeded portion of the charging energy must be removed after the required charging energy is obtained.

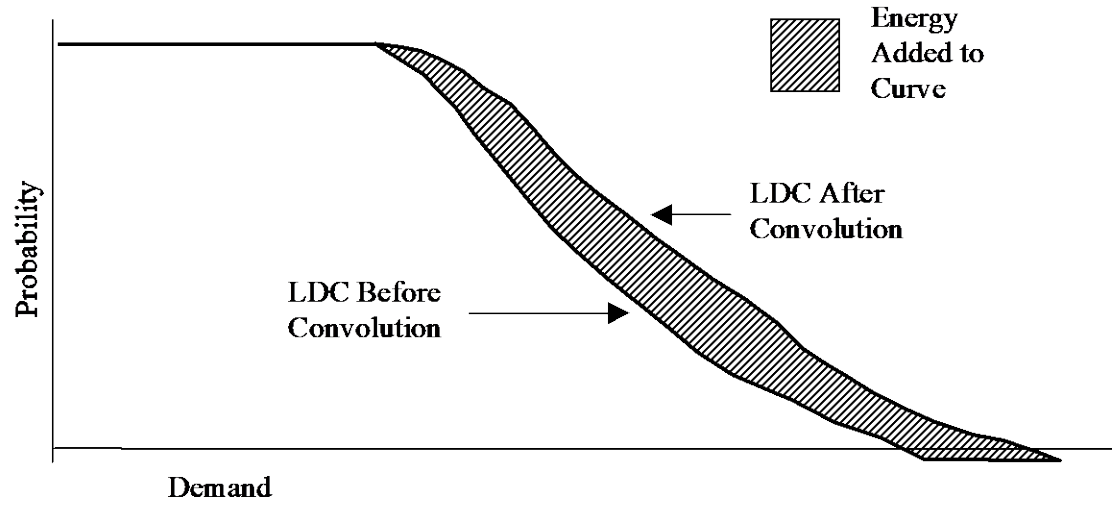

(a) Convolution of Charging Load onto LDC

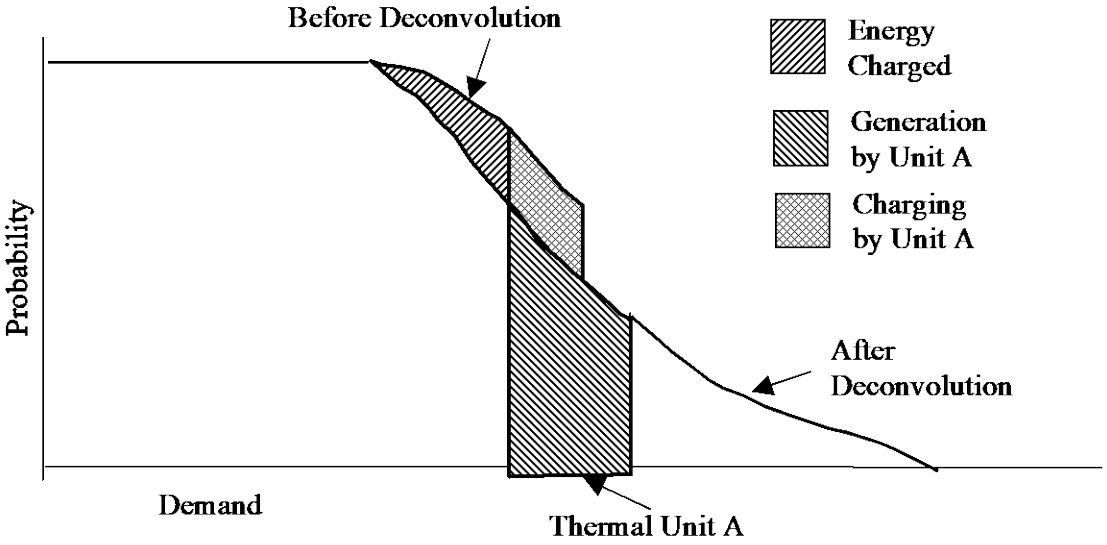

(b) Deconvolution of Charging Load from LDC

Figure 3-8
Dispatch of a Storage Charging Side

As EGEAS loads each succeeding thermal unit, it keeps track of how much energy has been charged – the area between the equivalent load duration curves (ELDC) with and without the charging load. When the required charging energy has been obtained, the fictitious LEP (which represents the storage charging side) is deconvolved so that no more energy is charged. The point where the required charging amount is achieved often falls in the middle of a thermal unit rather than at the boundary between units. Thus, the deconvolution of the fictitious LEP "splits" a thermal unit, as illustrated in the bottom portion (b) of Figure 3-8. Because of the energy invariance property of convolution, EGEAS does not need to determine the exact point for deconvolution. Instead, the energy generated by the last thermal unit providing charging energy

(Unit A in Figure 3-8) is adjusted by the algorithm to incorporate the correct amount of charging energy. Note that when a storage charging side is convolved or deconvolved, the loading point for the next generating unit (where it is dispatched) is not changed because the charging side provides no generation.

Two problems can occur when EGEAS models the dispatch of storage units with user-supplied generation values. The first problem is that the thermal units cannot provide sufficient charging energy. After all the thermal units have been dispatched, the fictitious LEP, which represents the storage charging load, is still convolved onto the ELDC because the required charging energy has not been found. Likely causes for this problem include an excess amount of storage capacity, a severe shortage of thermal capacity, and a relatively flat load distribution (high load factor). The first two causes occur most often in intermediate steps of an expansion optimization where unrealistic combinations of new units are sometimes generated. The last cause indicates that storage is not an appropriate technology because there are not enough off-peak load hours when a storage unit can be charged.

The second problem is that the storage operation is not economical. The marginal cost of providing the storage charging energy exceeds the marginal benefit of displacing peaking generation. Likely causes for this problem include an excess amount of storage capacity, a relatively high minimum load (base loaded units are fully used to meet load and cannot provide charging energy), and a relatively small difference in marginal operating cost between the units providing charging and the units being off-loaded. The first cause occurs most often in intermediate steps of an expansion optimization where unrealistic combinations of new units are sometimes generated; the last two causes indicate that the system does not need the amount of storage specified.

The basic solution to both problems is to reduce the storage generation. Thus, you should reduce the energy limit for the storage unit or ask EGEAS to optimize the storage operation with the medium-term decision. The medium-term optimization automatically avoids these problems.

Medium-Term Decision

The algorithm for determining the optimal amount of storage generation is based on two modes of operation: compulsory and economic. Compulsory operation occurs when thermal generation is not sufficient to serve the system demand and the storage unit can be used to reduce the expected unserved energy. An operation is considered economical when the thermal generation exceeds the system demand but the production cost can be reduced by operating the storage unit.

The algorithm has four main steps. In the first step, EGEAS determines the potential energy available for charging a storage project. In the second step, EGEAS determines the potential for displacing thermal peaking generation with storage generation. In the third step, EGEAS allocates the charging energy for compulsory operation to individual thermal units. In the fourth part, EGEAS determines the optimum amount of storage generation, allocates the required charging energy, and determines the energies off-loaded from peaking units. Following in outline form are the basic steps of the algorithm:

Step A - Charging Potential

A1. Perform production costing without any storage operation. Record the amount of energy generated by each thermal unit; this value represents the thermal unit's operation to meet the system load.

A2. Perform production costing again, with the storage charging load added to the initial LDC. (The charging side is convolved before any thermal units are dispatched and remains convolved while all the thermal units are loaded.) Record the amount of energy generated by each thermal unit; this value represents the thermal unit's operation to serve the system load plus the additional load imposed by the storage charging requirements.

A3. For each thermal unit, subtract the generation in Step A1 from the generation in Step A2 to obtain the potential energy available for charging.

Step B - Off-Loading Potential

B1. Perform production costing without any storage operation. Record the amount of energy generated by each thermal unit; this value represents the thermal unit's generation when not off-loaded (displaced) by storage operation. Note: this step is needed only when the storage charging and generating operations are separated into different time periods (see discussion at end of this subsection). Otherwise, the thermal generation values are taken from Step A1.

B2. As an extension to the production costing performed in Step B1 (or Step A1 if storage charging and generating operations are in the same period), load the storage generating side after all the thermal (and LEP) units. Record the amount of energy generated; this energy represents the potential for compulsory operation (reduction in unserved energy) by the storage project.

B3. Perform production costing again with the storage generating side loaded before any other units. Record the amount of energy generated by each thermal unit; this value represents the energy supplied to meet the system load, assuming the unit was displaced by the storage project.

B4. For each thermal unit, subtract the generation in Step B3 from the generation in Step B1 to obtain the potential off-loading (displacement) by the storage project.

Step C - Compulsory Operation

C1. Get the generation needed for compulsory operation from Step B2.

C2. Determine the charging energy needed for compulsory operation by dividing the compulsory generation by the storage efficiency.

C3. Sum the potential charging energies from Step A3 to obtain the total potential charging energy.

C4. If the total potential charging energy (Step C3) is less than or equal to the required charging energy (Step C2):

Set each thermal unit to produce the maximum possible charging energy and reduce the compulsory storage generation to reflect the available charging energy. Because all the charging energy is used for compulsory operation, economic operation is not possible.

C5. If the total potential charging energy (Step C3) is greater than the required charging energy (Step C2):

Set the compulsory storage generation to the Step C1 value. Allocate charging energy to the thermal units, up to their full potential (Step A3), in order of increasing marginal operating cost, until the required charging total is obtained. The remaining charging potential can be used for economic storage operation.

Step D - Economic Operation

D1. Select the thermal unit (P) with the highest marginal operating cost and a potential to be off-loaded by storage generation.

D2. Determine the marginal benefit for off-loading unit P:

$$MB = Z(P) - Z(S)$$

where: MB = marginal benefit ($/MWH)

Z(P) = marginal operating cost for unit P

Z(S) = marginal operating cost for the storage project

D3. Select the thermal unit (B) with the lowest marginal operating cost and a potential for charging the storage project.

D4. Determine the marginal cost for charging with unit B:

$$MC = Z(B)/Eff$$

where: MC = marginal cost ($/MWH)

Z(B) = marginal operating cost for unit B

Eff = storage efficiency

D5. If the marginal benefit (of off-loading unit P) exceeds the marginal cost (of charging with unit B), then economic operation is possible:

Increase the economic storage operation by the minimum of:

Potential charging from unit B, multiplied by the storage efficiency

Potential off-loading of unit P

Potential storage operation, determined by energy limit

Decrease each of the above quantities (potential charging by B, potential off-loading of P, potential storage operation) by the increased storage generation; at least one of these quantities goes to zero.

D6. Repeat Steps D3 through D5 until economic operation is no longer possible with unit P or the storage generation total reaches the energy limit.

D7. Repeat Steps D1 through D6 until no more units (P) can be economically off-loaded or the storage generation total reaches the energy limit.

EGEAS provides two options for allocating charging and off-loading energies to thermal units: deterministic and probabilistic. The algorithm outlined above pertains to deterministic allocation. With this option, charging energy is taken from the least expensive thermal unit up to its full

potential before charging energy from the next cheapest unit is considered. Similarly, energy is off-loaded from the most expensive unit up to its full potential before energy is off-loaded from the next most expensive unit. Thus, the deterministic option provides the most economic use of storage.

The probabilistic allocation option spreads charging and off-loading energies over a wider range of units. For compulsory operation, the charging energy is taken from all thermal units with charging potential. Energy is allocated to a particular unit in proportion to its charging potential. For example, if Unit A has twice the potential charging of Unit B, then Unit A provides twice as much compulsory operation as Unit B. Even the most expensive unit on the system may provide charging energy for compulsory operation. For economic operation, thermal units with potential charging energy are paired up in all combinations with thermal units with potential off-loaded energy. The storage operation resulting from each combination of units is dependent on the probabilities that the relevant units can provide charging and be off-loaded. Thus, the least expensive unit may not provide charging up to its full potential, even considering its contribution to off-loading all possible thermal units.

The rationale behind the probabilistic option is that storage dispatch based on a load duration curve (LDC) is often too optimistic. With the LDC methodology, the time sequence of events is lost. Thus, the level of storage operation developed with the deterministic option may be based in part on charging operations when the storage project is already fully charged and on generation when no stored energy remains. For many utility systems, the reduced storage operation determined with the probabilistic allocation option more closely matches actual system operation. However, each utility system is unique, so you must carefully choose the option that yields the best results.

There are three primary options for modeling storage:

No storage projects: You do not define storage projects. EGEAS operates each storage unit at the specified energy limit; i.e., EGEAS bypasses the medium-term optimization.

Storage projects at energy limit: You define storage projects and specify that projects should operate at their energy limits (i.e., EGEAS should bypass the medium-term optimization). You may specify deterministic or probabilistic allocation. With the deterministic option, the resulting storage operation is the same as when there are no projects (there may be minor differences when multiple storage units are combined in a project).

Storage projects with optimized energy use: You define storage projects and specify that EGEAS should determine storage operation with the medium-term optimization. You may specify deterministic or probabilistic allocation.

When the Benders' decomposition optimization method is selected, EGEAS always uses deterministic allocation (to satisfy conditions necessary for convergence to an optimum solution).

EGEAS can analyze storage operations for subyearly periods. If a year is broken down into multiple segments or subyears, EGEAS determines storage charging and generation levels for each of these periods. If the year, subyear, or segment is divided into subweeks, however, EGEAS places all storage generation in the peak subweek and all charging in the remaining subweek(s); the optimum storage use is optimized over the entire year, subyear, or segment, and

not for individual subweeks. Storage charging and generation are not allowed in the same subweek.

When there are multiple storage projects, EGEAS performs the calculations described above for each project. The main changes in the algorithm involve Steps A1, A2, B1, and B3. For the Nth project, EGEAS dispatches units in Step A1 to meet the system load plus charging load for the first N-1 projects. EGEAS dispatches units in Step A2 to meet the system load plus charging load for the first N projects. Thus, EGEAS reserves potential charging energies for the previous N-1 projects before calculating the potential charging energies for the Nth project. Similarly, for the production costing in Steps B1 and B3, EGEAS dispatches the generating sides for the first N-1 projects before any other units are loaded. Thus, the potential off-loading energies for project N assume that the thermal units have already been off-loaded for the first N-1 projects.

The order in which projects are modeled may be specified or you can request that EGEAS order the projects by decreasing energy limit, per unit of capacity. The order of the projects can affect the amount of energy generated by individual projects and also the total production cost, but the differences are usually small. In general, the algorithm is biased toward the projects modeled first.

Non-Dispatchable Technologies

Model Description

In EGEAS, non-dispatchable technologies (NDTs) represent sources of electric generation or load modification that are not controlled by the utility or that are strongly correlated with system demand. NDTs include: generation sources such as run-of-the-river hydro, wind, and solar; electricity purchases from co-generators, qualifying facilities, independent power producers; conservation programs based on insulation of homes and installation of energy-efficient appliances; and load management activities such as controlling water heaters and air conditioners. NDTs can also represent purchases from, or sales to, another utility, assuming the energy transfers have been scheduled for specific hours.

EGEAS makes a distinction between an NDT and an NDT resource. An NDT represents a specific hourly pattern of energy supply pertaining to a particular technology located at a particular location. For example, the hourly pattern would be different for solar and wind technologies; the pattern could also be different for two identical wind power stations installed at different locations; but, the hourly pattern would be the same for two solar photovoltaic units with the same design installed side by side. An NDT, reflecting the “energy source” availability, corresponds to the hourly output of one MW of installed capacity, assuming no equipment failures. Examples of this “energy source” availability include the solar radiation falling on a solar photovoltaic plant or the demand for steam from an industrial co-generator.

For a given NDT, there can be many NDT resources. Each of these resources has the same “energy source” availability and the same hourly output per unit of capacity. However, each resource can be unique in terms of capacity, forced outage rate, maintenance requirements, and costs. For example, two solar photovoltaic plants at the same site might have photovoltaic arrays of different sizes with varying hardware reliability and maintenance requirements.

For some NDTs, the utility is concerned with large numbers of very small units. For example, a wind farm can contain many wind generators and a load management program can involve control of many water heaters. For computational efficiency, EGEAS allows you to model

groups of identical NDT units. While the characteristics of the units are identical, their mechanical outages are treated probabilistically and are assumed to be independent. Thus, modeling accuracy would be sacrificed if EGEAS simply summed the capacities of all the units. Instead, EGEAS assumes a normal distribution for the total available capacity and replaces the individual units with a three-state representation of this distribution. For the three states, EGEAS uses the mean expected output (probability 0.5), one standard deviation above the mean (probability 0.25), and one standard deviation below the mean (probability 0.25).

For each NDT, you must input a utilization value for each hour of every year. Each utilization value gives the output per MW of capacity for a unit of that technology. For an energy sale, and for some load management activities, the load on the conventional generating units is increased in some hours; such increases in the load are indicated by negative utilization values. Thus, the utilization values should be between 0.0 and 1.0 to represent energy supplied to the system (or a reduction in the load), and between 0.0 and -1.0 to represent energy taken from the system (or an increase in the load). You can supply different hourly patterns for each year or use the same pattern over several years. Within a given year, you can supply typical daily patterns or separate values for every hour of the year (8736 values).

For each NDT resource, you must input values for capacity, forced outage rate, maintenance requirements, and costs. The capacity and forced outage rate can vary from year to year, but not seasonally. Basically, the same data is required for NDTs as other types of DSM and supply resources. EGEAS can model NDT units that have already been built (or are scheduled to be built) and can also optimize whether to build future NDT capacity. (In EGEAS terminology, an NDT unit can have an existing, committed, or generic status; see Chapter 2 for details.) When EGEAS models interconnections with another utility or a power pool, NDT units are allowed only on the primary system (System A).

Methodology

EGEAS models NDTs by modifying the initial system load. Load modification on an hourly basis would be prohibitively time-consuming (because EGEAS may need to consider many different combinations of NDT units), so EGEAS works with load duration curves. EGEAS treats the system demand and the output from each NDT as random variables and represents their dependencies by developing a joint probability distribution. The available capacity from each specific NDT unit is also a random variable; EGEAS assumes that the outages of individual units are independent. Thus, NDT generation is the product of two random variables: the energy source availability (which is interdependent with the load and other NDTs) and the generating capacity availability (which is independent of the other random variables).

The modified load distribution is given by:

$$Y^* = Y - \sum_i E_i \sum_j X_{ij}$$

where:

Y* = random variable representing the modified load
Y = random variable representing the initial system load
Ei = random variable representing the energy source available for technology i
Xij = random variable representing the generating capacity availability for unit j of technology i

Because the generating unit availabilities (X_{ij}) are statistically independent of the load (Y) and energy source availability (E_i) variables, EGEAS can determine the joint probability distribution of the (Y) and (E_i) variables just once and then use the results repeatedly to determine the modified load (Y^*) for different levels of NDT unit capacity. The probability distribution for the modified load is then converted into a load duration curve. EGEAS always uses the analytic (cumulants) representation for load duration curves when modeling NDTs.

EGEAS modifies the load for NDT output for each subweek of each year. Conventional units are dispatched at the annual, segment, or subweek level of detail. When the annual or segment level is selected, EGEAS aggregates the modified load over all the subweeks within a given period. The basic sequence of operations for load modification and production costing is:

1. Modify hourly loads for an entire year to match user-specified peak load and energy.
2. Modify hourly loads to account for fixed purchase and sale contracts (see the Purchase and Sale Contracts section). Note that these contracts are not modeled as NDT units.
3. Determine joint probability distribution for random variables representing the system load and the energy source availability for each NDT; this step is performed for each subweek.
4. Select new generating units, including NDT units, for modeling. (EGEAS may try many different combinations of units as part of the generation expansion optimization.)
5. Modify the load in each subweek to account for NDT generation.
6. Generate the modified LDC for each period by aggregating the subweek loads (if necessary).
7. Dispatch conventional generating units (thermal, hydro, storage) to meet the remaining system load.

For a conventional generating unit or DSM program, the marginal operating cost affects how the resource is dispatched. EGEAS dispatches resources with the smallest operating costs before more expensive resources to minimize the production cost. Thus, if the operating cost of a resource is increased, it tends to generate less energy. For an NDT resource, however, the marginal operating cost has no effect on the resource's output. The load modification process is independent of the NDT order and NDTs are modeled before EGEAS considers what conventional capacity is available. Thus, the output of an NDT is governed only by the hourly utilization values and the capacity installed; EGEAS takes the full NDT output, even when more economic alternatives are available. You can specify operating costs for NDT resources, however, and these costs are included in the total cost of an expansion plan. These costs also affect the generation expansion optimization if any of the NDT resources are candidates for future installation. EGEAS would tend to install NDT resources with low operating costs and conversely, tend not to install NDT resources with high operating costs.

Purchase and Sale Contracts

For many electric utilities, purchases and sales of energy are important facets of their generation planning. Energy purchases can help utilities through periods of low reserve margins until new capacity comes on-line. Energy sales can help utilities make better use of under-utilized resources. Basically, purchase and sale contracts help utilities match electricity supply to the energy demand.

EGEAS can model fixed committed contracts for sales and purchases. However, it cannot use these contracts as decision variables in the generation expansion optimization. To make decisions about entering into contracts, you can make comparative optimization runs with and without the potential contract. Another option is to model the contract as a non-dispatchable technology unit, which is allowed as a decision variable in the optimization. Many contracts can also be modeled as dispatchable resources, which can then be optimized. Note that firm contracts cannot be specified for interconnected systems (B through J). EGEAS uses either of two methods for modeling contracts.

The primary method involves modifications to the initial system load. Because the contracts are fixed, the load modification can be performed just one time. Thus, unlike the load modification for NDT units, the load modification for contracts is performed on an hourly basis, before the load duration curves are constructed. The primary method allows six types of contracts: peak, valley, and unscheduled purchases; and peak, valley, and unscheduled sales.

The second method involves the sale of excess energy from specified resources. After a resource has been dispatched, EGEAS determines whether the resource can generate any additional energy and treats such energy as a sale.

Input Requirements

For each firm contract, you must specify a name, the type of contract, starting and ending years, and four basic quantities that describe the essential characteristics of the contract. First, the capacity limit (MW) specifies an upper limit on the instantaneous magnitude of the purchase or sale. Second, the energy limit, specified as a GWH value, provides an upper limit on the energy transferred in a particular period. Alternatively, you may specify a capacity factor (percent) and EGEAS calculates the energy limit as the product of the capacity (converted to GW), capacity factor (converted to a decimal), and the number of hours in the period. Third, the capacity cost is given in dollars per megawatt per week ($/MW/week) based on the contract capacity limit. Fourth, the energy cost is given in dollars per megawatt hour ($/MWH) for the total amount of energy actually transferred (which may be less than the energy limit). Costs are in current year dollars (not base year dollars). For a sale contract, you supply capacity and energy revenues rather than costs. All four quantities can be changed from year to year and also from segment to segment within a given year.

An alternate option for the energy cost (or revenue) is to derive the $/MWH value from the operating cost of one or more resource units. Included in the operating cost are fuel costs, variable O&M costs, and any detailed costs that are modeled as variable O&M. You can specify up to six committed or existing resources; generic resources are not allowed. When determining the average, EGEAS weights resources by the amount of energy they supplied in the production costing dispatch. With this option, you may also specify a cost modifier (in $/MWH) to be added to the average resource cost.

You may also designate one or two preferred subweeks. For an excess energy type of contract, energy is sold only in the preferred subweek(s). For all other contract types, energy sales or purchases are allocated first to the preferred subweek(s); other subweeks are used only if the energy limit is not reached in the preferred periods. If no preferred subweeks are specified, EGEAS makes no distinction when allocating energy to the hours of a particular segment of the year.

The contract contribution to the system reserve margin is specified as a percentage of the contract capacity. If EGEAS calculates reserve margin on an annual basis, the percentage is applied to the average contract capacity over the entire year. If EGEAS calculates reserve margin on a subperiod basis, it uses the current contract capacity in each segment. For example, if a contract has a capacity of 100 MW in segments 2 and 3, but 0 MW in segments 1 and 4 (assuming 4 segments of 13 weeks each), the percentage is applied to 0, 100, 100, and 0 MW for reserve calculations in segments 1 through 4, respectively. However, for an annual reserve calculation, EGEAS applies the percentage to 50 MW, the average annual capacity. For a purchase, a positive value is treated as a contribution to the resources needed to meet the reserve margin requirements, while a negative value is treated as a reduction in the load. For a sale, a positive value is treated as a reduction in the resources available to meet the reserve requirements, while a negative value is treated as an increase in the load.

The spinning reserve contribution of a contract is specified as a percentage (zero or positive) of the contract capacity. The actual MW value of the contribution can change from year to year and from segment to segment as the contract's capacity changes. If the contract capacity varies from segment to segment but production costing is performed on an annual basis, the spinning reserve contribution is based on the average contract capacity for the entire year.

Methodology

The primary method for modeling contracts is to modify the initial hourly loads. The basic sequence of operations is:

1. Input the hourly load values for an entire year.
2. Apply a linear transformation to the hourly load values to achieve the proper annual peak load and energy.
3. Modify hourly loads to account for purchases and sales contracts, in the order in which the contracts are supplied.
4. Determine the load duration curves and perform orthogonalization with non-dispatchable technologies, and with the System B-J load, if any.

EGEAS analyzes contracts one week at a time. For a given segment, EGEAS uses the same capacity in every week and divides the segment energy equally among the weeks. For peak contracts, energy is allocated to hours in order of decreasing load. The hourly loads are decreased (for a purchase) or increased (for a sale) starting in the hour with the greatest load and continuing until the energy limit is reached. The load is decreased or increased by the contract capacity, except possibly in the last hour modified. If the load is modified in every hour before the energy limit is reached, then the energy purchased or sold is less than the energy limit for that week; the remaining energy is not transferred to another week. For valley contracts, the process starts with the hours having the smallest load. For unscheduled contracts, the loads are modified for hours chosen randomly. If one or more preferred subweeks are specified, loads are modified first in the designated subweek(s) before proceeding to the remaining subweek(s).

EGEAS has an option to generate an output file containing the net hourly loads after the contracts (except excess energy sales) have been applied. This file is in the same format as the hourly loads input file. Thus, the modified loads could be supplied as input to a later run. One application of this feature is the modification of interconnected system (B to J) loads to account for purchases and sales. EGEAS does not allow purchases and sales for interconnected systems,

but you can run EGEAS to preprocess the load for an interconnected system (treating it as if it were System A) before the analysis of the actual System A begins.

Excess energy sale contracts are not modeled in the same manner as other contracts. For these contracts, modeling begins after EGEAS has performed production costing to determine the generation by individual resources on the system. For each resource whose excess energy is to be sold, EGEAS compares the energy actually generated with the maximum energy the resource could produce. If the resource is not operating at its full capability, then EGEAS raises the resource's output up the maximum and sells the excess energy. If an energy limit (GWH) is supplied for the contract, then the excess energy sold cannot exceed that value. When subweeks are modeled and the potential sales for a particular segment exceed the energy limit, then the sales are reduced proportionately in all the subweeks of that segment. If a capacity factor limit (percent) is supplied for the contract, then that percentage of the potential energy is sold. EGEAS performs excess energy sale calculations at the same level of subperiod detail as used for production costing. If one or more preferred subweeks are specified, excess energy sales are limited to the preferred subweek(s) only.

For all types of contracts, EGEAS determines the capacity costs, based on the input contract capacity, for every segment of every year. Energy costs are based on the amount of energy actually sold or purchased. If energy costs are based on the average operating costs of one or more resources, EGEAS determines the energy costs after the production cost dispatch establishes the generation values to be used as weights in the averaging. EGEAS determines the net purchased energy cost by summing the capacity and energy costs over all purchase contracts and then subtracting the corresponding values (treated as income) for sales contracts.

EGEAS determines contract contributions to reserve margin and spinning reserve before optimizing or simulating expansion plans. Thus, reserve margin constraints reflect the contribution already provided by contracts, and spinning reserve requirements are reduced to reflect contract contributions before EGEAS starts modifying the dispatch order to meet spinning reserve requirements.

Maintenance Scheduling Options

In EGEAS, maintenance refers to planned periods when a unit is taken out of service for repairs or routine maintenance work. This is in contrast to forced outages, which may be anticipated but are not planned. For each unit, EGEAS allows two maintenance periods per year and each period must be an integral number of weeks in length. The maintenance requirements for a unit may vary from year to year. Maintenance is modeled by derating each generating unit's capacity by the fraction of time the unit is out of service. If the load duration curve is represented analytically (cumulants method), then the cumulants of a unit's outages are also derated. For example, if a unit has six weeks of maintenance, then the unit's capacity and all of its outage cumulants would be multiplied by the factor:

$$\frac{52-6}{52} = 0.8846$$

Maintenance Cycles

Because maintenance requirements may vary from year to year, EGEAS accepts maintenance data in the form of maintenance cycles. If a unit has a maintenance cycle of N years, then the

maintenance requirements for a particular year are repeated every N years. Up to two maintenance periods for a single year can be specified by entering two starting weeks. Because maintenance requirements may not settle into a regular pattern for several years after installation, the user can also specify several years of initial maintenance data before the cycle starts. For example, suppose a unit requires 12 weeks of maintenance in its first year of operation and 8 weeks of maintenance in its second year before settling down to a regular three-year cycle with 3, 4, and 6 weeks of maintenance. For this example, the user would specify 5 years of maintenance data (with 12, 8, 3, 4, and 6 weeks, respectively) and a 3-year cycle. Thus, the maintenance requirements for successive years would be:

12, 8 3, 4,6 3, 4,6 3, 4,6 3, 4,6...

For specifying maintenance requirements, the first year can be either:

the installation year of the unit,

the first year after the base year of the data base, or

the base year of the data base.

The installation year option is useful for generic units where the installation year is not known until the analysis is performed. The base year option is useful for existing units which may have been installed many years ago. With this option, you need not trace current maintenance cycles back to the installation year.

Maintenance Scheduling Algorithm

When EGEAS performs production costing on an annual basis, maintenance is modeled by derating each generating unit's capacity and cumulants for the fraction of time the unit is on maintenance in that year. EGEAS makes no attempt to allocate the unit's maintenance to any particular time during the year because all weeks are in the same period.

When EGEAS performs production costing on a subperiod basis, however, maintenance requirements can be scheduled either by the user or by the maintenance scheduler in EGEAS. You can schedule a unit's maintenance periods by specifying a starting week. EGEAS then derates the unit's capacity and cumulants in the appropriate segments of the year. As an example, assume the year is divided into 4 segments with 13 weeks each. If a unit has 6 weeks of maintenance starting in week 10, then that unit's maintenance will be split between segments 1 and 2. The first 4 weeks of maintenance (weeks 10 through 13) will be in segment 1, so the unit's capacity and cumulants will be multiplied by (13 - 4)/13. The remaining 2 weeks of maintenance will be in segment 2 (weeks 14 and 15), so the unit's capacity and cumulants will be multiplied by (13 - 2)/13 in segment 2.

The maintenance scheduler determines when to maintain units whose maintenance periods have not been pre-specified. EGEAS never overrides a starting week that you specify.

The objective of the maintenance scheduler is to levelize either the loss of load probability (LOLP) or the reserve margin over the segments of the year. Thus, EGEAS tends to schedule units for the spring and fall periods where the load, and consequently the reliability measures are higher/lower (whichever applies) while avoiding the peak summer and winter periods. The basic steps of the algorithm (assuming levelization of LOLP) are:

1. Determine the equivalent load duration curve (ELDC) for each segment of the year by convolving the outages of all generating units onto the initial load duration curve (LDC).

For units whose maintenance schedule is pre-specified, derate the capacity and cumulants in the appropriate segments; for all other units, ignore the maintenance requirements.

2. Calculate the LOLP for each segment of the year and rank the segments in order of increasing reliability.
3. Rank the units whose maintenance is to be scheduled in order of decreasing capacity.
4. Schedule the maintenance of the (remaining) unit with the highest capacity in the segment with the lowest LOLP.
5. If the maintenance period is too long (too many weeks) to fit in the chosen segment, extend the maintenance period into the segment immediately preceding or following the initially chosen segment. If both a preceding and following segment are available, choose the one with the lower LOLP. Extend the maintenance period to additional segments as necessary keeping the maintenance in contiguous weeks.
6. For each segment in which the unit is maintained, deconvolve the unit from the ELDC, derate the cumulants for the fraction of time the unit is maintained, and deconvolve the unit with the adjusted cumulants.
7. For each affected segment, recalculate the LOLP.
8. Repeat steps 4 through 7 until maintenance is scheduled for all units.

Notice that when a unit's maintenance is scheduled in a particular segment, the LOLP for that segment increases. EGEAS continues scheduling units for maintenance in that segment until it is no longer the most reliable. If you have chosen to have reserve margin levelized instead of LOLP, the same algorithm is used except that EGEAS schedules maintenance in the segment with the highest reserve margin instead of the segment with the lowest LOLP.

When segments are aggregated into subyearly periods, EGEAS uses the same algorithm with subyears in place of segments. Maintenance is never allocated to specific portions of a week (e.g., subweeks). Use of the maintenance scheduler is optional; when it is not used and a unit's maintenance period is not scheduled by the user, EGEAS spreads the maintenance requirements evenly over the entire year. For example, a unit with 6 weeks of maintenance has its capacity and cumulants multiplied by (52 - 6)/52 in every subperiod of the year.

There is a limitation on maintenance scheduling when it is used in conjunction with the Benders' decomposition analysis option. In most cases, the maintenance scheduler works with all the generating units: existing, committed, and generic. With the Benders' decomposition option, however, only existing and committed units can be scheduled; schedules for generic units can be pre-specified but cannot be determined by the maintenance scheduler. This restriction is necessary to guarantee generation of an optimal expansion plan. Maintenance schedule reports are available upon request.

Dual Fuels and Fuel Constraints

For thermal resources, you have the option of specifying either one or two fuel types. For hydro and storage resources, you can (optionally) specify one fuel type. Fuels can be shared by multiple resources or each resource can have its own unique fuel type. An upper limit can be placed on the available supply for any fuel type; a lower limit can be placed on the usage of one particular fuel type, to model a take-or-pay contract, such as for gas. For resources that have two fuels, you can pre-specify the percentage of a unit's generation to be powered with the secondary

fuel or you can let EGEAS set the percentage based on economics or available supply. EGEAS provides reports showing fuel usage for each resource and for each fuel type.

Input Requirements

Input requirements related to fuel usage fall into three main categories: input for generating units (or DSM programs), input for fuel types, and modeling options. For generating units, the key parameters are the heat rate, the fuel type(s), and for units with two fuels, the target, minimum, and maximum percentages for the secondary fuel. The heat rate, specified in units of BTU/kWh (or equivalently, MBTU/GWH), can be changed from year to year; it cannot be changed seasonally. For a unit with multiple blocks, a different heat rate can be given for each block to represent the unit's heat rate curve (see the section on Heat Rates in Chapter 2 for more details). If a unit has two fuels, the heat rate for the secondary fuel can be different from the one used for the primary fuel.

The target percentage for the secondary fuel specifies what percentage of a unit's generation is to be powered by the secondary fuel, the remainder being powered by the primary fuel. If not specified, the default is zero percent (for all primary fuel). EGEAS uses the target percentage if no fuel constraints are applied or if there is sufficient supply of the two fuels. The target percentage can be changed from year to year or seasonally. If desired, the percentage can be set to 100 percent for some parts of the year and 0 percent for others to reflect different fuel usage patterns in summer versus winter.

If fuel constraints are applied, EGEAS may need to change the percentage of a unit's generation that is powered by the secondary fuel. For such a situation, you can specify both minimum and maximum secondary fuel percentages. EGEAS will never violate these percentages, even curtailing the operation of a generating unit if necessary. As with the target percentage, the minimum and maximum percentages can be changed annually and seasonally.

For fuel types, the basic input parameters are the unit of mass (e.g. ton or barrel), the heat content (MBTU per unit of mass), and the cost ($/MBTU). The heat content of a fuel is a constant, but the cost can change both annually and seasonally. In addition, you can specify both upper and lower limits on the fuel usage. The upper limit represents the available fuel supply, while the lower limit represents a contractual lower limit on the amount of fuel which must be used (e.g., take-or-pay gas contract). While lower limits can be specified on any number of fuel types, EGEAS can impose a lower limit on only one fuel type in a particular analysis. Both lower and upper limits can be changed from year to year. Also, the limits can be allocated to individual segments of the year. If you do not pre-specify how the fuel limits are to be allocated, EGEAS will automatically allocate the fuel in proportion to the length of each segment; EGEAS does not optimize the allocation based on the load or where the fuel is most needed.

For modeling options, you can specify whether maximum and/or minimum fuel constraints are to be applied. When specified, maximum constraints are applied to all fuel types which have a limited supply. For a minimum constraint, you must also designate the single fuel type to which the constraint will apply. When fuel constraints are applied, EGEAS modifies the dispatch order, switches fuels, and/or curtails unit operation as needed to satisfy the constraints. In addition, you can specify fuel switching options for maximum and minimum fuel constraints, and also for emissions. The switching options apply only to units that have two fuels. If a switching option is turned on, EGEAS will change the percentage of a unit's generation that is powered by the secondary fuel, if necessary to satisfy the corresponding fuel or emission constraints. Switching

fuels for emissions applies only when the emissions production is lower with one fuel than with the other. If no switching options are turned on, EGEAS will keep each dual-fueled unit at the target percentage that you specify.

One additional input parameter is the dispatch order option, which specifies which costs are used in determining the dispatch order. As mentioned in the Dispatch Order section, there are five options:

Weighted average fuel cost plus variable O&M (default)

Weighted average fuel cost only

Variable O&M cost only

Primary fuel cost plus variable O&M

Primary fuel cost only

For the weighted average fuel cost, EGEAS uses the cost of the dual fuels, weighted according to the target secondary fuel percentage.

Methodology

The methodology used for modeling dual fuels and fuel constraints depends on the constraint options specified. One constant, however, is that EGEAS tries to minimize the production cost, subject to reliability, fuel, and other resource (e.g., emissions) constraints. Following are the algorithms EGEAS uses in six different cases:

Case 1. Maximum Fuel Constraints without Fuel Switching

Before starting the production costing dispatch, EGEAS determines the dispatch order. For a unit with two fuels, the program assumes that the target percentage of generation will be fueled with the secondary fuel. EGEAS also determines the available supply of each fuel type. During the dispatch, each unit is loaded on a trial basis and its fuel usage is determined. If the use of either primary or secondary fuel exceeds the available supply, the unit's generation is curtailed as much as necessary based on the fuel available. Because the fuel switching option is turned off, EGEAS will not switch fuels even if an abundant supply of one fuel could make up for the lack of the other fuel. Because the unit cannot generate the amount of energy required by a unit loaded at the current position under the LDC, EGEAS "delays" the unit – postponing its dispatch to a point where it can operate at full capacity and generate the amount of energy permitted by the available fuel. In essence, the unit is treated as a delayed hydro unit (see the earlier section on Limited Energy Plants).

After each unit is loaded the remaining fuel supply is updated to reflect the fuel used by that unit. If a unit's dispatch is delayed, the fuel supply is updated at that point rather than waiting until the unit is actually dispatched. In this way, the fuel is not used by some later unit, whose cost would be higher than the unit that was delayed. When a unit is delayed, its generation is determined, so the fuel is reserved for it.

Case 2. Maximum Fuel Constraints with Fuel Switching

As with Case 1, EGEAS begins by determining the dispatch order, assuming each dual-fueled unit is using its target secondary fuel percentage. EGEAS then dispatches the system, ignoring fuel constraints completely. After the dispatch is complete, the

program checks the amount of each fuel used. If the usage for each fuel type is less than or equal to the available supply, production costing is complete.

When at least one fuel constraint is not satisfied, however, EGEAS proceeds to the next step, which is fuel switching. For this step, EGEAS searches for dual-fueled units that use one of the fuel types whose usage is too high. If possible, the program changes the secondary fuel percentage to lower the use of the over-utilized fuel. To reduce primary fuel usage, the secondary fuel percentage is raised; to reduce the secondary fuel usage, the percentage is lowered. If both fuels are over-utilized, EGEAS cannot resolve the problem and leaves the fuel percentage unchanged. Note that when switching fuels, EGEAS will not violate the minimum or maximum secondary fuel percentages that you specify. You can specify whether the program should start fuel switching with the based-loaded units or with peaking units. Based on your input, the program proceeds through the units in dispatch order or reverse dispatch order. During fuel switching, the overall generation by a unit is not changed – only the amount generated with the primary and secondary fuels is changed. If at some point, the usage of all fuel types is within limits, the fuel switching process stops and the production costing is complete.

If fuel switching by itself cannot bring the usage of all fuels within limit, EGEAS proceeds to dispatch the system again. For this second dispatch, the program curtails the generation of individual units as needed to keep the fuel use within the limit (as in Case 1). The secondary fuel percentage for each unit is left at the value it had after fuel switching.

Case 3. Minimum Fuel Constraint without Fuel Switching

As with Case 1, EGEAS begins by determining the dispatch order, assuming each dual-fueled unit is using its target secondary fuel percentage. EGEAS then dispatches the system, ignoring fuel constraints completely. After the dispatch is complete, the program checks the fuel use for the designated fuel type. If the usage is at least the minimum required, production costing is complete.

When the fuel use is below the minimum, EGEAS dispatches the system again. This time, the program tries to increase the generation by units that use the designated fuel type. After the must-run units are dispatched, EGEAS searches for the next unit in the dispatch order that uses the designated fuel. The program dispatches this unit next and records how much more fuel is used than in the first dispatch. If the fuel use now exceeds the minimum, EGEAS delays the unit to where it will exactly use the desired amount of fuel. (To keep the production cost low, the program changes from the original (economic) dispatch order as little as possible). If the fuel use is still below the minimum, EGEAS searches for another unit that burns the designated fuel and repeats the process. If all units that burn the designated fuel are placed right after the must-run units and the minimum fuel use is still not obtained, EGEAS simply dispatches the rest of the units economically – there is no way to achieve the minimum fuel use. In this case, the program adds the cost of the unused portion of the fuel to the production cost; EGEAS assumes this is a take-or-pay situation where the utility must pay for the fuel even if it is not used.

Case 4. Minimum Fuel Constraint with Fuel Switching

As with Case 3, EGEAS begins by dispatching the system, ignoring the minimum fuel constraint. If the constraint is satisfied, production costing is complete. Otherwise, the program attempts to change the secondary fuel usage for fuel-fueled units that use the designated fuel type to increase its use. You can specify whether EGEAS should start switching with the base-loaded or peaking units. If the minimum fuel use can be achieved with fuel switching, production costing is complete. Otherwise, the program proceeds to dispatch the system again, as in Case 3. Units that burn the designated fuel type are moved earlier in the dispatch order to increase their generation, and in turn, their usage of the designated fuel type. If the minimum fuel use is not achieved, EGEAS adds the cost of the unused fuel to the production cost.

Case 5. Maximum and Minimum Fuel Constraints without Fuel Switching

If both maximum and minimum fuel constraints are requested with no fuel switching, EGEAS uses an algorithm that includes steps from Cases 1 and 3 above. The general process is as follows:

Step 1: Perform production costing. Curtail generation where necessary to satisfy maximum fuel constraints. Ignore minimum fuel constraint. If minimum fuel usage is met, process is complete.

Step 2: Perform production costing again, moving units that burn the designated fuel type (for minimum fuel constraint) right after the must-run units until the minimum fuel use is achieved. As in Step 1, curtail generation where necessary to satisfy the maximum fuel constraints.

Case 6. Maximum and Minimum Fuel Constrains with Fuel Switching

If both maximum and minimum fuel constraints are requested with fuel switching, EGEAS uses an algorithm that includes steps from Cases 2 and 4 above. The general process is as follows:

Step 1: Perform production costing, ignoring both maximum and minimum fuel constraints. If all constraints are satisfied, process is complete.

Step 2: Switch fuels, without changing generation totals, in attempt to satisfy fuel constraints. Switching is performed first for the minimum fuel constraint and then for the maximum fuel constraints. It is possible that switching for one constraint may reverse steps taken to satisfy another constraint. Because the maximum fuel constraints go last, they take precedence over the minimum fuel constraint. If all constraints are satisfied, process is complete.

Step 3: Perform production costing again. Curtail generation where necessary to satisfy maximum fuel constraints. Ignore minimum fuel constraint. If minimum fuel usage is met, process is complete.

Step 4: Perform production costing once more, moving units that burn the designated fuel type (for minimum fuel constraint) right after the must-run units until the minimum fuel use is achieved. As in Step 3, curtail generation where necessary to satisfy the maximum fuel constraints.

Representation of Load Duration Curve

For production costing, EGEAS uses load duration curves to represent the distribution of system demand. EGEAS gives you a choice of two methods for representing the initial and equivalent LDCs: piecewise linear and analytic. With the piecewise linear option, EGEAS represents the curve as a series of equally-spaced points, giving the height of the LDC. This option is available only with screening curves and for the calculation of incremental costs; it is not allowed for production costing. With the analytic option, EGEAS uses either a Gram-Charlier series or a multi-parameter Gamma distribution, based on either 6 or 8 moments of the load distribution. The two analytic methods are available for production costing and the calculation of incremental costs. Four operations must be performed on the LDC: convolution of unit outages, deconvolution of unit outages, calculation of the height of the LDC for some point, and calculation of the area under the curve between two points. The production costing algorithm described in earlier sections is the same with both representations; the only difference is in how these four basic operations are performed.

Piecewise Linear Representation

The piecewise linear representation of the LDC is based on a set of points, equally spaced along the demand axis. If there are N points, then the first point is at 1/N times the peak load and the last point is at the peak load. As units are convolved onto the LDC, the equivalent peak load shifts to the right by the capacity of the unit. EGEAS keeps the same spacing between points as on the initial LDC; the increased range of load values is accommodated by adding more points.

With the piecewise linear representation, EGEAS uses the trapezoidal rule to calculate areas (i.e., EGEAS approximates the area under the curve between two MW values by summing the areas of the trapezoids that comprise the desired area).

For a unit with n blocks, convolution of outages onto the LDC is given by the formula:

$$G'(L) = p * G(L) + \sum_{i=1}^{n} q_i * G(L - x_i)$$

where:

$G(L)$ = initial LDC at load level L

q_i = partial outage rate for block i

x_i = capacity for blocks i through n

p = availability for the entire unit (1 - sum of q_i)

$G'(L)$ = equivalent LDC after convolving unit

EGEAS uses this formula to determine the value of the new equivalent load duration curve (ELDC) for each of the equally spaced points (L). Because the capacities (x_i) are not usually a multiple of the spacing between points, the values for $G(L - x_i)$ are determined by linear interpolation.

The deconvolution formula for a multiple-block unit is:

$$G(L) = \frac{G'(L) - \sum_{i=1}^{n} q_i * G(L - x_i)}{p}$$

where G' is the ELDC before deconvolution and G is the ELDC after deconvolution. The deconvolution process is recursive because the value of G at point L depends on the values of G at lower load levels ($L - x_i$). This recursive property can make the deconvolution operation unstable because a small numerical error at one point can get propagated to higher load levels.

The accuracy of the piecewise linear method depends on the number of points used to represent the LDC. Accuracy improves as more points are used, but the computational burden increases as well. In addition, the piecewise linear method can be numerically unstable for large systems with repeated convolutions and deconvolutions.

Analytical Representation With Cumulants

For an analytical representation of the LDC, EGEAS uses either a Type A Gram-Charlier series or a multi-parameter Gamma distribution. Both are based on either 6 or 8 moments of the load distribution. A Gram-Charlier series has the desirable properties that it can be used to approximate any probability distribution and can be written in terms of the cumulants of that distribution. Cumulants, which are polynomial functions of the moments of a distribution, have the useful property that the cumulants of the sum of two independent random variables equal the sum of the independent variable cumulants. Thus, the convolution operation is reduced to mere addition.

To calculate the cumulants for the initial system load or for generating unit outages, EGEAS starts with the mean and the central moments of the distribution. Both the load and outage variables are discrete probability distributions, so the mean is given by:

$$m = \sum_{i=1}^{n} g_i x_i$$

where: n = number of discrete states for random variable x

x_i = value of x for state i

g_i = probability that variable is in state i

For the system demand:

n = number of hours in the year or subperiod

x_i = system demand for hour i

g_i = 1/N for all i

For a single-block resource unit:

n = 1

x_i = capacity for resource i

g_i = forced outage rate for resource i

For a multiple-block resource unit:

n = number of blocks

x_i = capacity summed over block i and all higher blocks

g_i = block outage rate for block i

The central moment (moment about the mean) of order r is given by:

$$c_r = \sum_{i=1}^{n} g_i (x_i - m)^r$$

Finally, the cumulants of the distribution are given by:

$$k_1 = m$$

$$k_2 = c_2$$

$$k_3 = c_3$$

$$k_4 = c_4 - 3c_2^2$$

$$k_5 = c_5 - 10 c_3 c_2$$

$$k_6 = c_6 - 15 c_4 c_2 - 10 c_3^2 + 30 c_2^3$$

$$k_7 = c_7 - 21 c_5 c_2 - 35 c_4 c_3 + 210 c_3 c_2^2$$

$$k_8 = c_8 - 28 c_6 c_2 - 56 c_5 c_3 - 35 c_4^2 + 420 c_4 c_2^2 + 560 c_3^2 c_2 - 630 c_2^4$$

Note that a cumulant of order r depends on moments of order less than or equal to order r. Thus, a series expansion that depends on r cumulants also depends on r moments.

The Type A Gram-Charlier series representation of the ELDC based on 6 cumulants is:

$$P6(x) = \int_x^{\infty} N(y)dy + \frac{G(3)}{3!} N^{(2)}(x) - \frac{G(4)}{4!} N^{(3)}(x) + \frac{G(5)}{5!} N^{(4)}(x) - \frac{G(6) + 10G^2(3)}{6!} N^{(5)}(x)$$

For 8 cumulants, the representation is:

$$P8(x) = P6(x) + \frac{G(7) + 35G(3)G(4)}{7!} N^{(6)}(x) + \frac{G(8) + 56G(3)G(5) + 35G^2(4)}{8!} N^{(7)}(x)$$

where:

$$N(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2} = \text{Normal Function}$$

$$N^{(r)}(x) = \frac{d^r}{dx^r} N(x) = (-1)^r H_r(x) N(x)$$

$H_r(x)$ = Hermite polynomial of order r

X = standardized equivalent load
(mean = 0.0 and standard deviation = 1.0)

$G(r) = k_r / (k_2)^{r/2}$ = Gram-Charlier coefficient of order r

k_r = rth order cumulant of the equivalent load
(system load plus unit outages)

When a resource is convolved onto an ELDC, EGEAS simply adds its cumulants to those of the ELDC to obtain the cumulants for the new ELDC. Similarly, when a generating unit is deconvolved, its cumulants are subtracted. Thus, the convolution and deconvolution operations are very efficient. For calculating the area under the ELDC, EGEAS uses the analytic

representation of the integral of P6 or P8. For example, the area under the tail of the ELDC from point t to infinity is given by:

$$F6(t) = \int_t^{\infty} P6(x)dx = \int_t^{\infty}\int_x^{\infty} N(y)dy\,dx - \frac{G(3)}{3!}N^{(1)}(t) + \frac{G(4)}{4!}N^{(2)}(t) - \frac{G(5)}{5!}N^{(3)}(t) + \frac{G(6)+10G^2(3)}{6!}N^{(4)}(t)$$

EGEAS evaluates the double integral of the normal function by linear interpolation from values in a table; it determines the derivatives of the Normal functions as products of Hermite polynomials and the Normal function itself. Note that all derivatives of the Normal function are zero at infinity.

For calculating the height of the ELDC, as needed to determine the LOLP, EGEAS could simply evaluate the Gram-Charlier series P6 or P8 at the desired demand level. However, the Gram-Charlier series based on only 6 or 8 cumulants does not provide a perfect fit of the ELDC, particularly in the tail of the curve. Area calculations, which cover a range of load levels, tend to be more accurate. Thus, EGEAS determines the height of the ELDC by computing the area under the curve in a small interval around the desired point and then dividing by the width of the interval. In other words, EGEAS uses the average height in a small interval rather than the height at a specific point.

The multi-parameter Gamma option uses the following 2n-parameter distribution:

$$f(x) = \frac{1}{\Gamma(\alpha)}\sum_{i=1}^{n} \omega_i\,\rho_i e^{-\rho_i x}(\rho_i x)^{\alpha-1}$$

where ρ_i = scale parameter of the *i*-th gamma distribution (*i* = 1,n)

α = shape parameter (same for all the distributions)

ω_i = weighting factor for the *i*-th gamma distribution (*i* = 1,n-1)

() = gamma function

The 2n parameters are determined by matching the moments of the multi-parameter gamma distribution to those of the equivalent load duration curve (derived from the cumulants).

Comparison of Representations

Both the piecewise linear and analytic representations of the LDC produce reasonably accurate cost and reliability values in the context of long-range generation planning. However, they differ in both accuracy and efficiency.

The piecewise linear method usually produces a better fit to the initial load duration curve. With this better fit, production costing results are more accurate, particularly the energy generated by peaking units, the LOLP, and the expected unserved energy -- quantities which depend on the values of the ELDC near the tail of the curve. However, for large systems, with many convolutions and deconvolutions, the instability of the algorithm may cause small errors to compound that will negate the method's inherent accuracy advantage. Also, the piecewise linear method is slower than the analytic method. This is a serious disadvantage when used for generation expansion optimization, which requires many production costing calculations. Thus, the piecewise linear method is allowed only with the screening curve method and for incremental cost calculations.; it is not allowed for production costing.

The two analytical methods tend to be less accurate than the piecewise linear method, but the accuracy is usually adequate for long-range planning. In particular, the generation by base and intermediate units and the total production cost are usually very close to those obtained with the piecewise linear method. However, calculations out at the tail of the ELDC, which are needed for determining peaking generation, LOLP, and unserved energy, are less accurate. The analytic methods are stable and relatively immune to propagation of errors with large numbers of convolutions and deconvolutions. Also, the methods are very fast, making them the natural choice for large optimization problems.

As for the two analytical methods, the multi-parameter gamma distribution provides better accuracy than the Gram-Charlier series in determining LOLP, unserved energy, and peaking unit capacity factors. The improved accuracy can lead to better maintenance scheduling and modeling of NDT units. The gamma method, however, increases run times significantly, particularly with a Benders decomposition algorithm. EGEAS does provide a "mixed cumulants" option for user with Benders, which uses both methods. The gamma distribution is used for production costing and reliability, but the Gram-Charlier series is used for calculating dual multipliers (derivatives). This option offers a compromise, providing better accuracy than the Gram-Charlier option, but much faster execution time than the multi-parameters Gamma option.

4
OPTIMIZATION ANALYSIS

Overview

EGEAS is designed to determine optimum expansion plans, or simulate pre-specified plans, for an electric utility. The dollar calculations within EGEAS are based on engineering economic principles. You may choose between two objective functions: present worth of revenue requirements or levelized system average rate. The optimization can include both supply-side and demand-side options. By supplying the necessary costs and choosing the appropriate objective function, you can simulate various DSM tests, such as the total resource cost test, the rate impact measure, the no-losers test, and the societal test. EGEAS provides four tools for analyzing expansion plans: screening curves, Benders' decomposition, dynamic programming, and pre-specified pathway. These tools are described in later sections and in the following chapters.

In addition to the above mentioned economic objective functions, EGEAS also provides a financial screening capability which can alert system planners to potential financing and ratemaking problems with a particular expansion plan. Such problems may arise when capital-intensive supply options are being pursued. The cost analysis capabilities of EGEAS complement the expansion optimization; they are not a substitute for comprehensive financial simulations. The level of detail is limited to reduce computation time and to avoid burdensome data requirements.

While the cost analysis capabilities of EGEAS provide useful information about the rate and financial implications of implementing a specific expansion plan, detailed financial simulation models can provide more accurate assessments. In some cases, a utility can circumvent or at least alleviate the problems signaled by EGEAS' economic indicators by implementing alternative ratemaking and financing policies. Thus, comprehensive financial and economic analyses are needed to validate final capacity expansion plans. While the cost analysis features are designed primarily for investor-owned utilities, they provide much useful information for municipal and cooperative utilities as well.

EGEAS provides three economic feasibility measures: interest coverage ratio, earning assets ratio, and system average rate. The interest coverage ratio, which gives the ratio of earnings before interest and taxes to the interest payments on fixed-income obligations, is widely used in the electric utility industry and the financial community. The earning assets ratio, which gives the fraction of assets that are in rate base (and thus providing a return to the utility), can be valuable for municipal and cooperative utilities, whose ratemaking and financial policies often differ from those of investor-owned utilities. The system average cost, which gives the total revenue requirements per MWH of energy sales, is useful as a market or rate indicator for any utility. EGEAS specifically looks at percentage increases in system average cost from year to year.

EGEAS can calculate and report these economic feasibility measures for each year of the study period. In addition, EGEAS can determine multiple-year averages. For the interest coverage ratio and the earning assets ratio, the multiple-year averages are not true averages, but are close. For

the system average rate, EGEAS calculates both single-year increases and compound average increases over several years.

In addition to calculating economic feasibility measures, EGEAS can impose constraints on these measures during generation expansion optimization. Such constraints can help you find the least cost expansion plan, in engineering economic terms, that is also feasible from a financing standpoint. With the generalized Benders' decomposition method, EGEAS can impose constraints on interest coverage ratios and earning assets ratios. With the dynamic programming method, constraints cannot be imposed during the optimization process. Instead, EGEAS evaluates the optimal and suboptimal plans in a cost postprocessor. Plans are marked feasible or infeasible based on the values of the economic feasibility measures. You can then decide which plans to investigate in more detail.

Objective Functions

EGEAS allows you to choose one of two objective functions for the expansion optimization: present worth of revenue requirements or levelized system average rate. However, the system rate option is allowed only with the dynamic programming method. Included in the revenue requirements are the following costs:

Capital fixed charges for installing new generating units and starting new DSM programs; capital fixed charges for existing and committed resources are optional.

Capital expenses for major plant upgrades or other capital improvements, for resources that have already been installed.

Operating costs for all resources, including fuel costs, fixed O&M costs, and variable O&M costs.

Detailed costs, such as DSM customer costs, rebound benefits associated with DSM programs, customer benefits from DSM, transmission/distribution costs, and other costs.

Emission allowance costs.

Cost of purchased energy contracts.

Cost of energy purchases from other systems with multi-area modeling.

Unserved energy costs.

The following revenues and credits are also included in the objective function:

Emission allowance credits

Revenues from energy sale contracts

Revenues from energy sales to other systems with multi-area modeling

The inclusion of many of these costs and credits depends on input options that you specify. In particular, there are options for each type of detailed costs, for emissions costs and credits, and for multi-area modeling. Also, some costs are excluded with certain optimization methods (e.g., detailed costs are not allowed with Benders' decomposition). Costs are summed over all years of the study period, plus the extension period, if one is specified; all costs are discounted to the base year.

For the levelized system average rate, EGEAS first calculates the system average rate in each year as the total system cost divided by the system sales. Included in the system cost are all the

costs listed above for the revenue requirements, except that some detailed costs (e.g., DSM customer costs) can be excluded. EGEAS also allows you to specify additional capital and operating costs (associated with "non-EGEAS assets") for the rate calculation. Examples of such costs are capital fixed charges for existing units (if not supplied for each individual unit), transmission and distribution costs, and general or administrative costs. System sales are determined from the system load by subtracting the contributions of DSM programs and accounting for the line losses. The levelized system average rate is determined from the present worth sum of the annual rates by using the following formula:

$$SARLEV = \left(\sum_{y=1}^{N} \frac{SAR(y)}{(1+R)^{y}} \right) * \left(\frac{R*(1+R)^{N}}{(1+R)^{N}-1} \right)$$

where: SARLEV = levelized system average rate

SAR(y) = system average rate for year y

R = discount rate, decimal

N = number of years

y = year, where 1 = first study period year

The summation is taken over all years of the study period, and optionally, the extension period as well. Refer to Chapter 7 for details on the calculation of the system average rate.

Screening Tests for DSM

Least-cost utility planning usually involves one or more screening tests used to determine the cost-effectiveness of demand-side management alternatives. By supplying the necessary costs and selecting the appropriate objective function, you can simulate several different screening tests. Following is a list of several DSM tests and how they can be simulated with EGEAS.

Total Resource Cost Test (TRC)

The TRC test evaluates whether the program cost per unit of the electricity saved is less than the cost per unit of the power supply avoided. The program benefits are the utility's avoided energy and capacity costs, while the program costs include utility and customer costs for the DSM program. The present worth of revenue requirements objective function should be specified for this test.

Rate Impact Measure

This test, sometimes termed the "rate impact measure/nonparticipant's test," analyzes the extent to which unit rates for electricity and gas are altered by the implementation of a DSM program. It is a measure of the change in revenue requirements, usually expressed on a "per sales" basis.

A DSM program will not increase rates if, after the implementation of the program, the present value revenue requirements per kilowatt-hour (kWh) are equal to or less than the present value revenue requirements per kWh before implementation of the program. The system average rate objective function should be specified for this test.

No-Losers Test

Under the no-losers test, if the total cost of the DSM program is equal to or less than the utility's avoided cost then the DSM program is cost effective. The test incorporates decreases or increases in sales revenues due to DSM program so that the impacts on nonparticipants can be measured. The no-losers test also measures DSM program effects on participants, nonparticipants, and the utility. A program passes the test only if none of those groups lose as a result of program implementation.

Societal Test

The societal test evaluates the total cost of the DSM program to society and determines if these costs are less than the total value of the resources saved by avoiding energy production.

Costs include all resources used to produce energy savings. Benefits include all resources saved by avoiding production of electricity, including environmental benefits, and any other social benefits external to the transaction between the utility and the customers. The present worth of revenue requirements objective function should be specified for this test.

Decision Variables

The decision variables for the expansion optimization are the numbers of planning alternatives of each type to be installed in each study period year. Each planning alternative can represent one of the following actions:

Installation of a new resource (generating unit or DSM program)

Installation of a new resource to operate for one year only (single-year alternative)

Retirement of one or more (up to 4) existing resources

Combined installation of a new resource and retirement of existing resources

No action (dummy alternative)

The installation of one alternative can be made dependent on the installation of one or more other alternatives, so that EGEAS can model the following situations:

Installation of a new resource

Implementation of a short-term measure (e.g., purchases) for a single year

Replacement of refurbishment of an existing resource

Phased installation of a new resource over several years, optionally at pre-determined intervals

Retirement of an existing resource

Choice of one from among several mutually-exclusive options.

The installation of a new resource is allowed with all analysis options; the other types of planning alternatives are allowed only with the dynamic program or pre-specified pathway. Refer to the Decision Variables and Dependency Constraints sections of Chapter 6 for more details on the different types of decision variables.

For a given analysis, you determine which types of resources should be considered and you may limit any type to a specific range of study years. For the screening curve and Benders' decomposition methods, EGEAS uses continuous real variables. For example, the optimum expansion plan might include 2.37 units of one type and 0.08 units of another type. For the

dynamic programming method, decision variables are discrete integer variables. For example, the optimum expansion plan might include 3 units of one type and 1 unit of another type, but never 2.37 units. For the pre-specified pathway, you can choose to simulate plans with either fractional or integer numbers of new resources.

Constraints

EGEAS imposes constraints on the generation expansion optimization to restrict the analysis to expansion plans that are feasible in terms of system reliability, resource utilization, technology development, financing capability, and regulatory approvals. Constraints may also reflect artificial limits imposed to promote computational efficiency. EGEAS gives you considerable flexibility in selecting which constraints to impose, when to impose the constraints, and what the limiting values should be (in most cases, the limits can be changed from year to year). In a few instances, EGEAS requires that specific constraints be imposed. Some constraint types are not allowed with all the optimization methods. Following are the types of constraints EGEAS can impose:

Reliability Constraints

Reserve Margin: The minimum and maximum reserve margin in a given study year or in each subperiod of a year. The user-specified limits, which may vary from year to year, are given as percentages of the peak load. Reserve margin is discussed in the **Reliability Measures** section in Chapter 2.

Loss of Load Hours (LOLH): The minimum and maximum LOLH in a given study year or in each subperiod of a year. The user-specified limits, which may vary from year to year, are given as hours per year. LOLH is discussed in the **Reliability Measures** section in Chapter 2.

Unserved energy: Upper limit on the unserved energy in a given study year or in each subperiod of a year. The user-specified limits, which may vary from year to year, are given as percentages of the initial system energy. Unserved energy is defined in the **Reliability Measures** section in Chapter 2.

Tunnel Constraints

Alternative Limitation Constraints: Lower and upper limits on the cumulative number of units of a particular planning alternative that can be installed up through a given study year. Alternative limitation constraints are not allowed with single-year alternatives. You may supply different limits for each alternative in each study year.

Alternative Installation Constraints: Lower and upper limits on the number of units of a particular planning alternative installed in a given study year. Alternative installation constraints are allowed both with conventional (multi-year) and single-year alternatives.

Financial Constraints

Earning assets ratio: Lower limit on the earning assets ratio for the utility in a single year, or on the average value of the earning assets ratio over several years. You may specify different limits for each year and different limits for single-year and multiple-year constraints. The earning assets ratio is defined in the **Feasibility Measures** section of Chapter 7.

Interest coverage ratio: Lower limit on the interest coverage ratio for the utility in a single year, or on the average value of the interest coverage ratio over several years. The user may specify

different limits for each year and different limits for single-year and multiple-year constraints. The interest coverage ratio is defined in the **Feasibility Measures** section of Chapter 7.

Percentage increase in system average rate: Upper limit on the percentage increase in the system average rate over the previous year value, or on the compound average increase over several years. The user may specify different limits for each year and different limits for single-year and multiple-year increases. System average rate is defined in the **Feasibility Measures** section of Chapter 7.

Fuel Constraints

Fuel availability: Upper limit on the amount of fuel, of a particular type, consumed in a given year. Fuel limits can be changed from year to year and the fuel supply can also be allocated to specific segments in a year. Fuel usage is described in Chapter 3.

Minimum fuel usage: Lower limit on the amount of fuel, of a particular type, that must be used in a given year. The limit can be changed from year to year and can be allocated to specific segments in a year. A minimum of fuel usage constraints can be used to model a take-or-pay fuel contract; a minimum fuel usage constraint can only be specified for one fuel type. Fuel usage is described in Chapter 3.

Emission Constraints

Unit emissions: Upper limit on the amount of emissions of a particular type, from a unit in a given year. Unit emission limits can be changed from year to year and annual limits can be allocated to specific segments in a year. See Chapter 9 for a complete description of the environmental capabilities of EGEAS.

Site emissions: Upper limit on the total amount of emissions of a particular type, from a group of units, in a given year. Site emission limits cannot be changed from year to year; the annual limit is always allocated to segments in proportion to the lengths of the segments. See Chapter 9 for a complete description of the environmental capabilities of EGEAS.

System emissions: Upper limit on the total amount of emissions of a particular type, from all the units on a system, in a given year. System emission limits can be changed from year to year and annual limits can be allocated to specific segments in a year. See Chapter 9 for a complete description of the environmental capabilities of EGEAS.

Emission rate reduction: Upper limit on the rate of emissions (e.g. tons CO_2 per GWH) produced by a particular group of units. Constraints can be satisfied by installing qualifying new "clean" generating units to the group of polluting units.

Other Constraints

Renewable Portfolio Standard (RPS): Lower limit on the percent of generation provided by renewable technology (e.g. solar, wind) generating units.

Optimization Methods

EGEAS provides four tools for analyzing expansion plans: screening curves, generalized Benders' decomposition, dynamic program (DP), and pre-specified pathway simulation. Each tool has unique strengths and weaknesses. While some planning problems can be solved using one of the tools, other problems can be solved most efficiently with a combination of tools.

The analysis options in EGEAS fall into a natural progression from simple screening methods (screening curves) to sophisticated optimization (Benders and DP) and simulation (pathway) methods. In broad terms, the simple methods take very little computer time to execute but use somewhat crude approximations; they are suited to quick comparisons of large numbers of planning alternatives. The sophisticated methods provide greater modeling accuracy and detail, but require substantially more computer time; they are suited to detailed optimization with a limited number of alternatives. The pathway simulation allows detailed evaluation of a small number of predetermined expansion plans.

For a generation expansion study, you may wish to use several of the methods. At the early stages, you can select the screening curve to compare a large number of planning alternatives. Alternatives that are clearly not competitive can be eliminated before performing more detailed (and costly) optimization. With a reduced set of alternatives, you can select the Benders or DP optimization to determine an optimal expansion plan; with the DP, you can also obtain a set of suboptimal plans. The pathway can be used to simulate the optimal or suboptimal plans with more detail than is computationally feasible during optimization, or the pathway can be used to evaluate minor variations in the optimal plan.

Following are some of the major advantages and disadvantages of each of the expansion analysis tools.

Screening Curves: The screening curve method executes very quickly and can handle large numbers of planning alternatives; the data requirements are minimal. However, the method is crude because it ignores existing and committed resources; it cannot handle LEP, storage, or NDT resources; it does not consider outages or maintenance requirements; and it cannot model subperiods, fuel and emission constraints, or interconnections to another utility.

Benders' Decomposition: Benders' decomposition can handle a moderate number of planning alternatives, can model all types of resources, performs probabilistic production costing, can apply economic feasibility constraints, and can determine the marginal effect of constraints on the optimization. However, it requires a moderate amount of execution time, installs fractional numbers of resources, and cannot model must-run units, spinning reserve, fuel and emission constraints, or interconnections to another utility.

Dynamic Program: The dynamic program can model all types of resources, performs detailed production costing with must-run resources and spinning reserve, can model interconnections to another utility, can generate suboptimal as well as the optimal expansion plan, and can evaluate the plans in terms of economic feasibility. However, it requires a considerable amount of execution time even for a limited number of planning alternatives, it cannot model annual changes in the characteristics of planning alternatives, and it cannot model alternatives with short operating lives.

Pre-specified Pathway: The pathway can simulate a predetermined expansion plan in great detail, but it does not perform generation expansion optimization.

Table 4-1 shows the primary modeling capabilities of these optimization tools.

Table 4-1
Comparison of Optimization Methods

	Screening Curve	Benders Decomposition	Dynamic Program
Computer execution time	Short	Moderate	Long
Number of planning alternatives (for reasonable execution time)	50	20	15
Types of units			
Thermal	Yes	Yes	Yes
Limited energy plant	No	Yes	Yes
Storage	No	Yes	Yes
Non-dispatchable technology	No	Yes	Yes
Types of planning alternatives			
New unit installation	Yes	Yes	Yes
Existing unit retirement	No	No	Yes
Phased installation	No	No	Yes
Major plant upgrade	No	No	Yes
Fuel Switching	No	No	Yes
Single-year options	No	No	Yes
Mutually-exclusive alternatives	No	No	Yes
Proper modeling of changes in characteristics based on unit age	No	Yes	Yes
Committed/existing units	Yes	Yes	Yes
Planning alternatives	No	Yes	No
Probabilistic dispatch	No	Yes	Yes
Probabilistic forced outages	No	Yes	Yes
Multiple loading blocks	No	No	Yes
Must run units and spinning reserve	No	No	Yes
Representation of LDC			
Piecewise linear	Yes	No	Yes
Analytic (cumulants)	No	Yes	Yes
Subperiods	No	Yes	Yes

Table 4-1 (continued)
Comparison of Optimization Methods

	Screening Curve	Benders Decomposition	Dynamic Program
Maintenance scheduling			
Existing and committed units	No	Yes	Yes
Planning alternatives	No	No	Yes
End effects modeling	Yes	Yes	Yes
Purchase and sale contracts	Yes	Yes	Yes
Interconnections for reserve sharing and economy interchange	No	No	Yes
Automatic avoided costs analysis	No	Yes	Yes
Constraints			
Reserve margin	No	Yes	Yes
Loss of load probability	No	No	Yes
Expected unserved energy	No	Yes	Yes
Tunnel	No	Yes	Yes
Fuel usage	No	No	Yes
Emissions	No	No	Yes
Interest coverage ratio	No	Yes	Yes
Earning assets ratio	No	Yes	Yes
Increase in system rate	No	No	Yes

Screening Curves

The screening curve method provides a quick comparison of planning alternates based on fixed and operating costs. It was designed primarily for analyzing the last year of a study period, but can be used for other years. The method is fairly crude, but is useful for preliminary analyses.

Although the screening analysis is intended primarily for thermal alternatives, non-thermal alternatives (hydro, storage, and NDT) are included in the cost tables for comparison. The energy limits for the non-thermal alternatives are ignored, as are the costs associated with charging storage resources. For an alternative that does not involve installation of a new resource (e.g., a dummy alternative or an alternative that merely forces an existing resource to retire), the costs are set to zero. If an alternative involves both the retirement of an existing resource and the installation of a new resource, only the new resource installation is considered. Screening curve analysis ignores any limitations imposed by dependency constraints (i.e., prerequisite alternatives) and by mutually-exclusive alternative constraints. Single-year alternatives are treated the same as multi-year alternatives.

When analyzing alternatives for a particular year, EGEAS assumes that all the alternatives are installed in that year. If an alternative is not available for installation in that year, it is excluded. Included in the analysis are costs over the period from the given year to the end of the study period, or to the end of the extension period if one has been specified. If an alternative retires before the end of this period, the program assumes that it is replaced with an identical resource. If a replacement resource retires in this period, the program replaces this resource also. Refer to the section on End Effects in Chapter 2 for more detail on the modeling of replacement resources.

Screening curve analysis involves three stages: calculating the costs for each alternative, determining the best alternative at each capacity factor, and selecting resources to meet the system load. For the first stage, EGEAS determines the total cost of each planning alternative as a function of capacity factor. The fixed portion of the costs includes capital fixed charges, fixed O&M costs, and any detailed costs (e.g., customer costs) that are modeled as capital or fixed O&M costs. The variable portion of the costs includes fuel, variable O&M, and any detailed costs that are modeled as variable O&M costs. EGEAS calculates the costs in terms of dollars per kilowatt of capacity. (The program uses the rated capacity of the resource, ignoring any annual or seasonal changes). Both the fixed and variable costs are discounted to the base year and summed over all years from the analysis year to the end of the extension period.

EGEAS assumes that the curve giving total cost as a function of capacity factor is a straight line, as shown in the top half of Figure 4-1. Each line on the graph represents a different planning alternative. The intercept on the vertical (cost) axis is the fixed cost. The slope is the marginal operating cost. In Figure 4-1, alternative 1 has the lowest fixed cost but the highest marginal operating cost, while alternative 5 has a high fixed cost but the lowest marginal operating cost.

To illustrate how EGEAS calculates the costs used for screening curve analysis, consider the following simple example:

Base year	1994
Discount rate	8.0%
Study period	1994 to 2002
Extension period	none
Planning alternative	
Rated capacity	100 MW
Heat rate	10,000 Btu/kWh
Fuel cost	1.0 $/MBtu (1994 $)
Fixed O&M cost	20.0 $/kW (1994 $)
Variable O&M cost	5.0 $/MWh (1994 $)
Capital fixed charges (levelized)	100.0 $/kW
Escalation rate for fuel, O&M costs	5.0%
Year of analysis	2000

The fixed portion of the costs would be:

Capital cost (for years 2000 to 2002, discounted to the base year 1994):

$$100/(1.08)^6 + 100/(1.08)^7 + 100/(1.08)^8 = 175.39\,\$/\text{kW}$$

Fixed O&M cost (escalated by 5% from 1994 to years 2000, 2001, 2002, and the discounted back to 1994):

$$20*(1.05/1.08)^6 + 20*(1.05)/1.08)^7 + 20*(1.05/1.08)^8 = 49.28\,\$/\text{kW}$$

Total fixed cost:

$$175.39 + 49.28 = 224.67\ \$/kW$$

The variable portion of the costs (assuming a capacity factor of 100 percent) would be:

Generation for one kilowatt of capacity for one year:

$$(1\,kW) * (8{,}736\ hours) = 8{,}736\ kWh = 8.736\ MWh$$

Variable O&M cost (escalated by 5% from 1994 to years 2000, 2001, 2002, and then discounted back to 1994):

$$(8.736\ MWh) * (5\$/MWh) / (1\,kW) = 43.68\$/kW$$

$$(43.68) * (1.05/1.08)^6 + (43.68) * (1.05/1.08)^7 + (43.68) * (1.05/1.08)^8 = 107.62\$/kW$$

Fuel cost (generation times heat rate times fuel cost, escalated and discounted):

$$(8.736\ MWh) * (10{,}000\ Btu/kWh) * (1\$/MBtu) / (1\,kW) * (1{,}000\ kWh/MWh) / (10^6\ Btu/MBtu) = 87.36\$/kW$$

$$(87.36\$/kW) * (1.05/1.08)^6 + (87.36\$/kW) * (1.05/1.08)^7 + (87.36\$/kW) * (1.05/1.08)^8 = 215.23\$/kW$$

Total variable cost:

$$107.62 + 215.23 = 322.85\$/kW$$

The total cost per kilowatt of capacity, based on capacity factor is:

Capacity Factor	Cost ($/kW)
0%	224.67
10	259.96
20	289.24
30	321.52
40	353.81
50	386.10
60	418.38
70	450.66
80	482.95
90	515.24
100	547.52

For the second stage of analysis, EGEAS determines the best (least expensive) alternative for each capacity factor. Only thermal alternatives are considered for this stage. For each alternative that is the best at some capacity factor, EGEAS shows the range over which that alternative is best. For example, one alternative might be the best alternative for capacity factors 18.6 to 35.4 percent. For the case shown in Figure 4-1, alternative 1 has the lowest cost for small capacity factors, alternatives 3 and 4 have the lowest cost for intermediate capacity factors, and alternative 5 has the lowest cost for high capacity factors. Thus, if the system needed peaking capacity, alternative 1 would be a good choice because of its low fixed cost. If the system needed baseloaded capacity, alternative 5 would be a good choice because of its low marginal operating cost. Alternatives 3 and 4 would be good choices for intermediate (cycling) resources. In this

example, alternative 2 is never the best alternative and thus is not a good choice for any type of operation.

For the third and final stage of analysis, EGEAS determines how many units of each resource to install to meet the system load. Alternatives are loaded deterministically (ignoring maintenance and forced outages) under the load duration curve (LDC) as shown in the bottom half of Figure 4-1. the alternative loaded at each location is the one with the lowest cost for operating at a capacity factor matching the percent of time on the LDC. In this example shown, alternative 1 is used as the peaking alternative, alternatives 3 and 4 are loaded for intermediate (cycling) operation, and alternative 5 is loaded for base-load operation. The vertical height of the band for each alternative determines the optimum amount of capacity for that alternative.

In general, the screening curve underestimates the amount of capacity needed because outages are ignored. On the other hand, no provision is made for existing or committed units, so the amount of new capacity may be overestimated. The main purpose of this third stage of analysis is to show, in broad terms, the ideal mix of resources well into the future where the existing system has essentially been retired.

Pre-Specified Pathway

With the pre-specified pathway option, EGEAS simulates an expansion plan specified by you. You define the plan by supplying the number of units of each planning alternative to be installed in each study period year. You have the option of specifying either integer or fractional numbers of resources. Multiple pathway plans can be specified to be calculated in a single EGEAS run.

For the integer option, EGEAS uses the last step of the dynamic programming method to simulate the expansion plan. For the fractional option, EGEAS uses the last step of the Benders' decomposition method. The capabilities and limitations of the pathway are the same as for the relevant optimization method. Note, however, that estimating installation years is no problem because the pathway knows the installation dates of all the planning alternative resources.

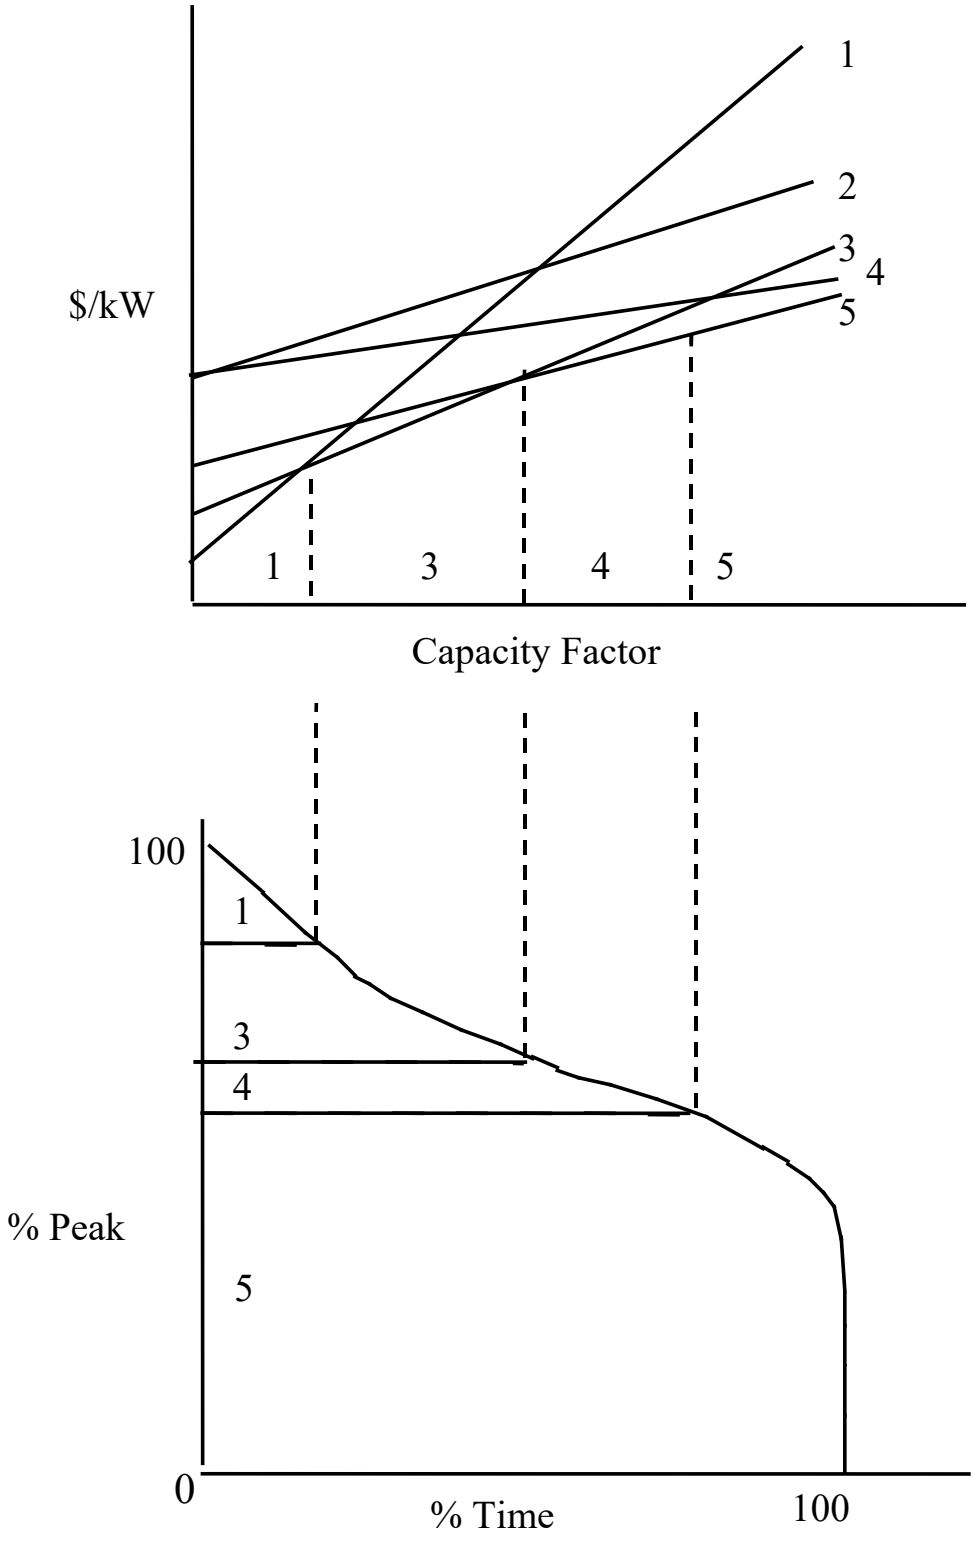

Figure 4-1
Screening Curves

5
BENDERS' DECOMPOSITION

Overview

The generalized Benders' decomposition method is based on the iterative use of a linear master problem and a non-linear subproblem. The master problem selects the capacity additions for trial expansion plans and the subproblem simulates the trial plans. The probabilistic production costing in the subproblem provides accurate plan simulation while the linearity of the master problem provides computational efficiency in the optimization.

Objective Function

For Benders' decomposition, the only objective function is present worth of revenue requirements. Included are capital and operating costs for all resources and unserved energy costs. Benders does not include detailed costs and emission allowance costs. The capital costs and fixed O&M costs for existing and committed units are predetermined constants; even though they are included in the objective function, they have no effect on the optimization.

For optimization, Benders requires that the unserved energy be priced at least as high as the most expensive resource. Thus, Benders determines its own price for unserved energy. For a given year, Benders sets the unserved energy price ($/MWH) to the marginal operating cost of the most expensive resource. For a storage resource, Benders includes the cost of the charging energy assuming that such charging is supplied by the most expensive thermal resource.

As described more fully in the **Methodology** section, Benders always applies an unserved energy constraint to limit the unserved energy to a specified percent of the initial energy. Because Benders assigns a high price to unserved energy, it may be cost effective to reduce the unserved energy below the allowed value. Nevertheless, Benders ignores such benefits; during optimization, the Benders subproblem stops dispatching resources when the unserved energy declines to the constraint limit.

After the optimization is complete, Benders simulates the optimal expansion plan to generate detailed information for reports. For this simulation, Benders makes several changes in the expansion plan cost. First, Benders dispatches all resources; it does not stop when the unserved energy limit is reached. Second, Benders calculates the unserved energy cost based on the input price rather than the internally-generated price used during optimization; the input price may be lower or higher than the internally-generated value. Third, Benders includes the costs associated with purchase and sale contracts. Benders does not model interconnections for economy interchange, so economy purchase costs are not included.

Decision Variables

The decision variables for Benders are continuous real variables representing the number of units of each capacity planning alternative installed in each study period year. Benders can model thermal, limited energy, storage, and non-dispatchable technology alternatives. Benders cannot model planning alternatives that represent retirement or replacement of existing units, and it

cannot consider dependencies among alternatives (e.g., phased installation or mutually-exclusive options).

Constraints

With Benders, EGEAS can impose the following types of constraints:

Reserve margin: Lower and upper limits on the reserve margin for a given study year. You can request minimum constraints only or both minimum and maximum constraints; constraints are applied in every study year or not at all. The limits can be changed from year to year.

Unserved energy: Upper limit on the expected unserved energy for a given study year. If Benders performs production costing for subyearly time periods, it imposes a limit on unserved energy in each subperiod; the same limit (percentage of initial energy) applies in every subperiod, but the limit can be changed from year to year. Note that Benders always imposes unserved energy constraints; you cannot turn them off.

Tunnel: Two types of tunnel constraints are available. Lower and upper limits on the cumulative number of resources of a particular planning alternative installed up through a given study year, and single year limits. The single-year limits can be useful in preventing EGEAS from deriving plans which have installations bunched in certain years. Thus, the resulting expansion plans are more realistic from a financing point of view.

- Earning assets ratio: Lower limit on the earning assets ratio for a given study year or group or years. Benders applies the constraints in every study year or not at all. Because constraints must be linear, Benders transforms a ratio constraint by multiplying by the denominator. Thus, the constraint is changed from:

  $$\frac{\text{Earning Assets}}{\text{Total Assets}} > \text{Limit}$$

  to:

  $$\text{Limit} * \text{Total Assets} < \text{Earning Assets}$$

  Because both the earning assets and the total assets are linear functions of the number of planning alternatives installed, Benders can convert the second form into a linear constraint for the master problem.

- Interest coverage ratio: Lower limit on the interest coverage ratio for a given study year or group of years. Benders applies the constraints in every study year or not at all. As with the earning assets ratio constraints, interest coverage ratio constraints are converted into a linear form.

See the description of constraints in Chapter 4 for more information about constraints.

Methodology

With the generalized Benders' decomposition methodology, the generation expansion optimization problem is decoupled into a linear master problem and a non-linear subproblem. Benders alternates between the two problems, with each passing new information to the other. The master problem determines a trial expansion plan based on a linear formulation of the optimization problem. If the problem were indeed linear, the optimization would be complete

with no need for further iterations. However, the problem is not linear, so the master problem sends the trial expansion plan to the subproblem for evaluation. The subproblem simulates the trial expansion plan and generates additional linear constraints to better approximate the non-linear cost and unserved energy functions. The subproblem sends these constraints back to the master problem for the next iteration. The iterative process continues until the method converges to an optimum solution. Figure 5-1 shows the relationship between the master problem and the subproblem.

Cost Constraints

The objective function for the linear master problem is the total expansion plan cost, which includes capital, operating and unserved energy costs. The capital cost for existing and committed resources is a constant, while the capital cost for planning alternatives is a linear function of the numbers of new resources installed. The operating costs can either increase or decrease as more resources are installed. Increases result from additional generation by the new resources (reductions in unserved energy). Decreases result from the displacement of generation from old inefficient resources by generation from new efficient resources. The unserved energy cost always decreases with the addition of new resources. The overall total of operating and unserved energy costs decreases because the cost of unserved energy is greater than the cost of supplying the energy. The operating and unserved energy costs are not linear functions of the number of resources installed because resources are dispatched only as needed to meet the system load and each succeeding resource sees a smaller remaining load (unserved energy) to be met. As more and more resources are installed, the operating and unserved energy costs approach constant values because additional resources have smaller and smaller effects.

For more details on production costing, refer to Chapter 3. Thus, the total expansion plan cost is not a linear function of the decision variables, the number of new resources installed. Figure 5-2 illustrates the behavior of the major cost components for a problem with just one planning alternative and one study year. In the figure, the magnitude of each cost component is the distance between two adjacent curves.

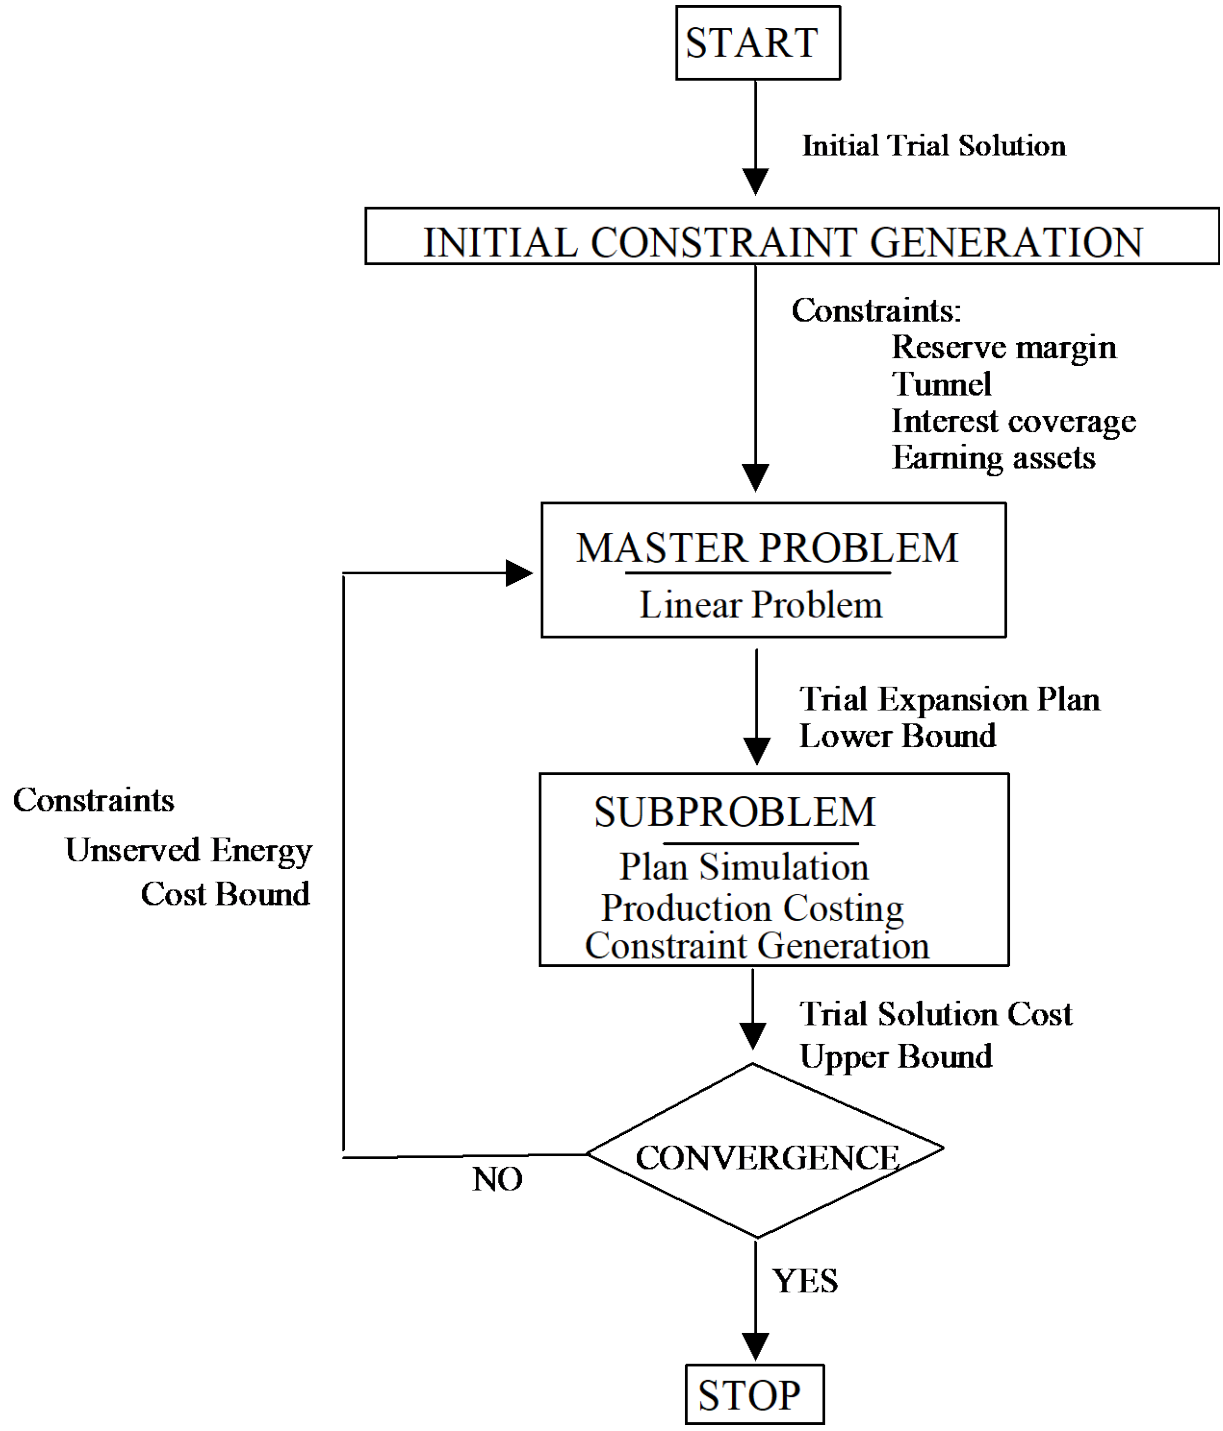

Figure 5-1
Benders Master Problem and Subproblem

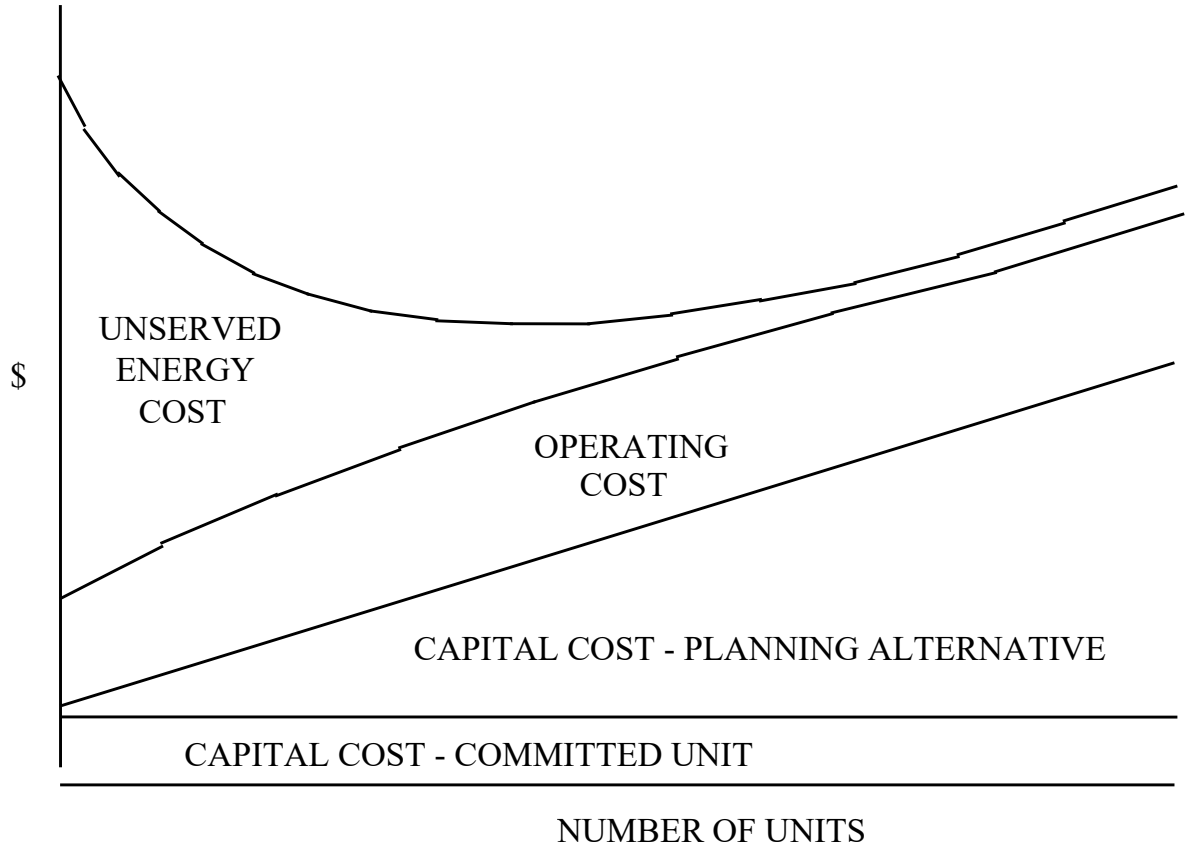

Figure 5-2
Expansion Plant Cost Components

To represent the non-linear cost in the linear master problem, Benders adds a new variable (Z) to the linear problem; the objective is to minimize Z. A series of linear cost constraints approximates the non-linear relationship between Z and the number of planning alternatives installed. At each iteration of the Benders optimization, the subproblem calculates the cost for a given trial solution and also the derivatives of cost with respect to the decision variables (these derivatives are also referred to as dual multipliers). The subproblem uses these quantities to construct a constraint that approximates Z in the region of the current trial solution. Essentially, this constraint represents a tangent plane to the non-linear cost surface.

Figure 5-3 shows how the non-linear cost function is approximated with a piecewise linear set of constraints for a problem with one planning alternative and one study year. For the first iteration (top graph), the number of resources is x_1. The subproblem determines the expansion plan cost and the slope of the tangent to the cost at x_1. Because the cost function is concave upward, the true cost must be above and to the right of the tangent line, so the tangent line forms a linear constraint on the variable Z. By itself, this constraint provides a poor approximation to the non-linear cost function, but the approximation improves as more constraints are added. The master problem uses this constraint in determining the next trial solution, x_2. Because the objective is to minimize Z, the master problem will choose a point to the right of where the cost constraint intersects the horizontal axis. While not shown in Figure 5-3, other types of constraints (e.g., reserve margin, tunnel, unserved energy) also affect the choice of x_2.

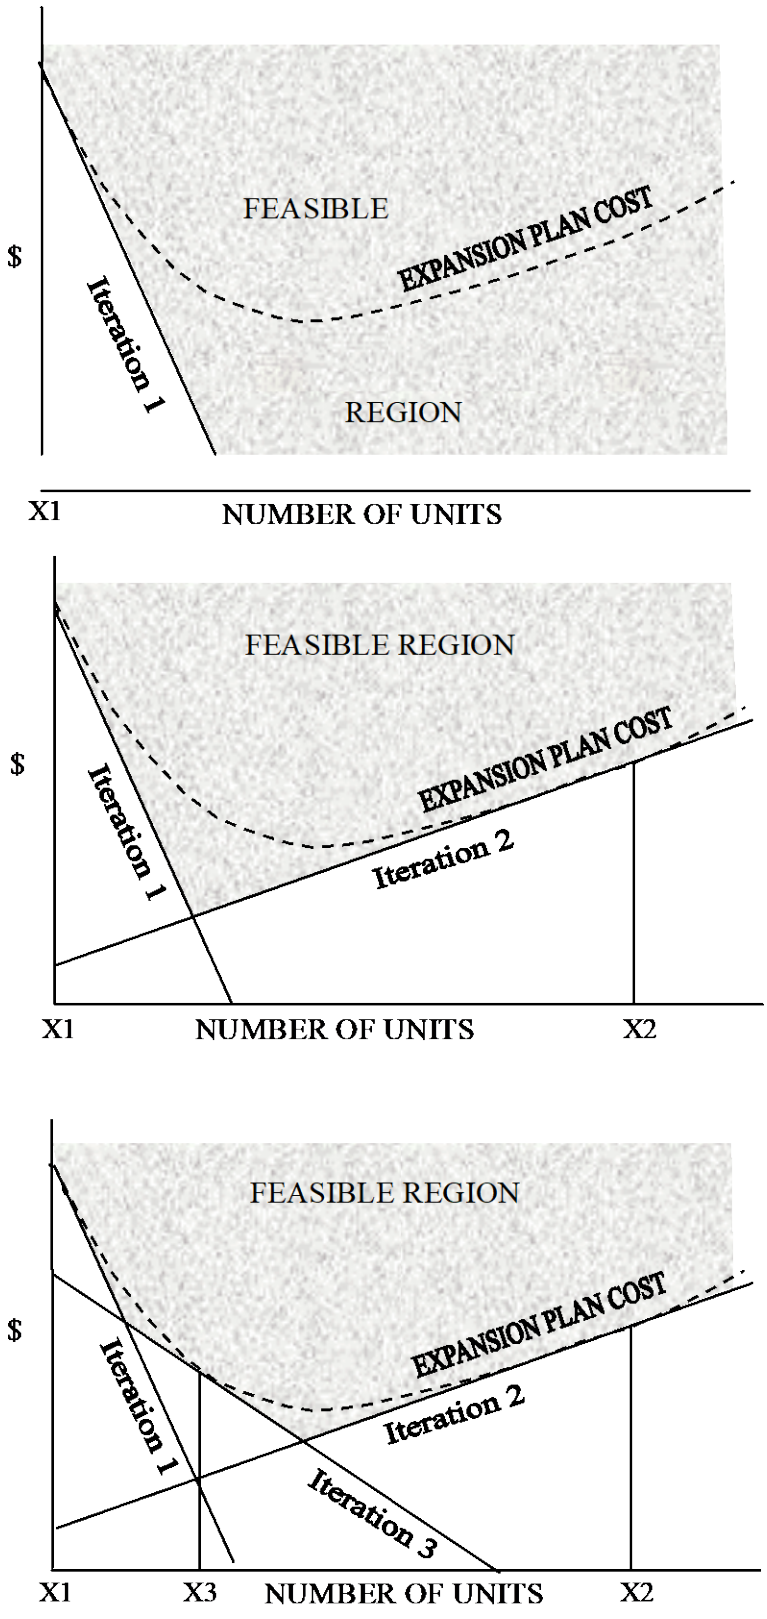

Figure 5-3
Benders Cost Constraints

For the second iteration (middle graph in Figure 5-3), the number of resources is x_2. At this point the subproblem determines the tangent to the cost function and uses this tangent as the cost constraint for the second iteration. Thus, the true cost function lies in the V-shaped region above the two cost constraint lines. In the absence of other binding constraints, the master problem chooses the intersection of these two lines as the next trial solution, x_3.

For the third iteration (bottom graph in Figure 5-3), the number of units is x_3. As with the first two points, the subproblem constructs the tangent to the cost function at this point and uses it as a cost constraint. The master problem then chooses the point in the feasible region, above the cost constraints, with the smallest cost; this point is where the iteration 2 and iteration 3 constraints intersect. With each new constraint, Benders has a better approximation to the non-linear cost constraint. The process continues until the trial solution is sufficiently close to the optimum value. The test for convergence is discussed in a later section.

Unserved Energy Constraints

Benders imposes an upper limit on the amount of unserved energy allowed in each study period year. Unserved energy, however, is not a linear function of the number of new resources installed, as illustrated in Figure 5-4. Thus, Benders relies on a series of linear approximations to the non-linear unserved energy function in much the same manner as it approximates the cost function with linear cost constraints. The unserved energy function is always concave upward because each additional resource that is included has a smaller effect on the unserved energy than the preceding unit. The unserved energy approaches zero as more resources are loaded. For production costing during optimization, however, the subproblem always stops dispatching resources when the unserved energy reaches the constraint limit. Thus, during optimization, the unserved energy never falls below the constraint limit.

When the subproblem evaluates a trial solution, it determines the amount of unserved energy in each study year and the slope of the tangent to the unserved energy curve. Figure 5-4 shows the tangent to the curve at a trial point x. The subproblem uses this tangent as a linear approximation to the unserved energy function. Based on this approximation, the subproblem determines the number of resources required for satisfying the unserved energy constraint. In Figure 5-4, the approximated unserved energy is lower than the upper limit when more than A resources are installed. Thus, the subproblem creates an unserved energy constraint that forces the master problem to install at least A resources in the next trial expansion plan. In reality, more than A resources are needed to satisfy the unserved energy constraint, but the subproblems for later iterations will generate constraints at points closer to the correct value (B). For multi-dimensional problems, the subproblem uses a tangent hyperplane to the unserved energy surface and the constraint is a linear function of all the planning alternative types.

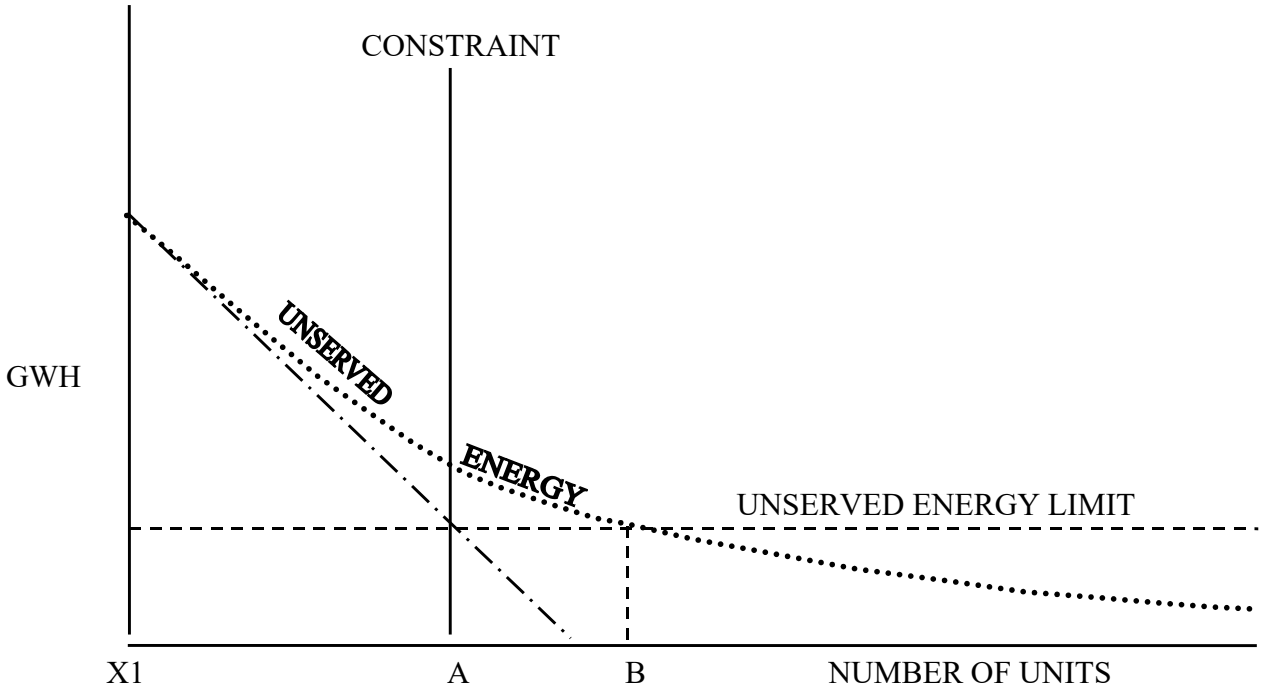

Figure 5-4
Benders Unserved Energy Constraint

For each iteration of the Benders optimization, the subproblem evaluates a different expansion plan and generates a new set of unserved energy constraints based on linear approximations to the unserved energy functions at that point. The subproblem generates a constraint for each study period year in which the unserved energy constraint is not satisfied; if the unserved energy is less than the upper limit, no constraint is generated.

If subperiods are modeled, the subproblem calculates the unserved energy for each subperiod of the year. If the unserved energy exceeds the upper limit in any subperiod, then a constraint is generated for that year. This constraint aggregates information from every subperiod in which the unserved energy is too high; it ignores subperiods where the unserved energy is within the allowed limit. For example, if the unserved energy exceeded the limit in subperiods 2 and 4 of a year with 6 subperiods, then the subproblem would generate an annual constraint with contributions from subperiods 2 and 4; it would ignore subperiods 1, 3, 5, and 6. On the other hand, if the unserved energy were within the limit for all 6 subperiods, then no constraint would be generated for that year. For each variable, the coefficient for the aggregate constraint is the sum of the corresponding values in the chosen subperiods; the right hand side of the aggregate constraint is the sum of the right hand sides. By aggregating the subperiod constraints, Benders reduces the size of the master problem to be solved. While applying an aggregated constraint is not as strict a requirement as applying individual subperiod constraints, the result is the same because the more reliable subperiods drop out of the aggregation as the iterative process approaches a feasible solution.

Subproblem

For each iteration of the optimization process, the Benders subproblem simulates the current trial solution (expansion plan). This simulation includes calculation of expected unserved energy values, productions costs, capital costs, and total expansion plan costs. Subperiod modeling

detail is allowed for production costing. The subproblem also generates unserved energy and cost constraints.

To generate constraints, the subproblem first calculates a set of dual multipliers. These dual multipliers are actually derivatives with respect to the number of new resources of a particular planning alternative. The dual multiplier for unserved energy in a particular year indicates how the unserved energy changes when the number of resources changes. (For a one-dimension problem as illustrated in Figure 5-4, the dual multiplier is the slope of the tangent to the unserved energy function.) The subproblem uses these dual multipliers with the current unserved energy value to construct the unserved energy constraint for the relevant year.

In a similar manner, the subproblem calculates dual multipliers for the production cost. After discounting these dual multipliers to the base year (the units are dollars per unit installed), the subproblem sums the multipliers over all the years, including the extension period. The resulting number represents how the production cost changes when the number of resources changes. The subproblem then adds the capital and fixed O&M costs for a single resource to obtain the rate of change for the total expansion plan cost with respect to the number of resources installed. (For a one-dimensional problem as illustrated in Figure 5-3, this combined rate of change is the slope of the tangent to the cost curve.) Thus, the subproblem can construct a cost constraint based on the current trial solution.

One additional function of the subproblem is the calculation of an upper bound on the cost of the optimal expansion plan. By definition, the cost of the optimal expansion plan must not exceed the cost of any feasible plan (a plan that satisfies all the constraints). Thus, the cost of any feasible plan is itself an upper bound on the cost of the optimal plan. If the current trial solution satisfies the unserved energy constraint in every time period, then the plan is feasible and its cost is an upper bound. If the current plan is not feasible, the subproblem converts it into a feasible plan by adding more resources of one of the planning alternatives. (The specific alternative used is chosen arbitrarily; any alternative can be used.) The subproblem adds just enough capacity to exactly satisfy the unserved energy constraints. By adding the capital and operating costs of these additional resources to the cost of the infeasible plan, the subproblem obtains the cost of a feasible plan, which then is used as an upper bound.

Master Problem

At each iteration of the Benders optimization, the master problem solves a linear problem. The objective function is the variable Z (described in the **Cost Constraints** section), which represents the total expansion plan cost. The constraints, including those listed in the **Constraints** section and the cost constraints that bound the cost variable Z, can be divided into two classes: initial constraints and subproblem constraints.

Initial constraints place bounds on quantities that are linear functions of the decision variables. This class includes reserve margin, tunnel, earning assets ratio, and interest coverage ratio constraints. Benders generates the constraints in this class just once and then includes them in the master problem optimization at every iteration.

Subproblem constraints impose limits on quantities that are not linear functions of the decision variables. This class includes the unserved energy and cost constraints. At each iteration, the subproblem constructs a constraint based on a linear approximation of the constrained quantity in the region of the current trial solution. Thus, the subproblem keeps generating new constraints

for the master problem. Note that the constraints from earlier iterations are retained, so the size of the master problem becomes larger for each succeeding iteration. For each iteration, one cost constraint encompasses the entire study and extension periods, but a separate unserved energy constraint applies for each year in which the unserved energy exceeds the allowed value.

Any linear problem can be constructed in either primal or dual form. So far, this discussion has focused on the primal form. The dual form has one variable for each primal constraint and one constraint for each primal variable. Thus, when a problem is converted from primal to dual form, the variables and constraints are interchanged. In the primal form of the Benders master problem, new constraints are added at each iteration; in the dual form, new variables are added. Because adding new variables to an existing problem is easier than adding new constraints, Benders formulates the master problem in its dual form. Thus, for example, there are dual variables for the reserve margin, tunnel, and cost constraints, and there are dual constraints associated with the numbers of planning alternatives installed.

For solving the linear problem, Benders usually uses the dual simplex method; however, you can specify that Benders should use the primal simplex method instead. In its dual form, the Benders master problem is very degenerate, so the primal simplex method sometimes requires many internal operations to find a solution. In most cases, the dual simplex method is better suited to solving the master problem in the dual form. For more information about linear problems, the reader should consult a textbook on optimization methods.

From the linear problem, the master problem extracts three types of useful information. First, the values of the decision variables (numbers of new resources to install) define the current trial expansion plan. With the master problem in its dual form, the numbers of new resources are obtained from the dual variables. Second, the value of the objective function serves as a lower limit on the cost of the optimal expansion plan because the cost constraints form lower bounds for the non-linear cost function (see Figure 5-3). Third, the dual multipliers provide shadow prices for the constraints. With the master problem in its dual form, the shadow prices are obtained from the primal variables. Note that these dual multipliers should be distinguished from the dual multipliers generated by the subproblem (see preceding subsection). For a particular constraint, the dual multiplier gives the rate of change in the cost of the optimum solution with respect to the value of the constraint bound. Thus, for example, if the dual multiplier for a reserve margin constraint is 2.0, then the cost of the optimal solution would decline at the rate of 2.0 million (base year) dollars for every MW by which the reserve margin constraint was relaxed. The units for the dual multiplier are millions of dollars per unit in which the constraint bound is measured (e.g., MW, GWH, resources installed). For cost constraints, the dual multipliers are dimensionless (actually millions of dollars per millions of dollars).

Convergence

At the end of every iteration, Benders checks whether the optimization has converged to the optimal expansion plan. For testing convergence, Benders uses the lower bound on the cost of the optimum plan generated by the master problem and the upper bound generated by the subproblem. The lower bound increases from iteration to iteration, so each iteration provides a better estimate of the optimal cost. The upper bound can be higher or lower than the value for a previous iteration, but the basic trend is downward. Thus, Benders keeps track of the lower upper bound, which is the minimum of all upper bounds up through the current iteration. Benders assumes convergence is achieved when the current lower bound and the lower upper bound are

sufficiently close. You can specify the convergence criterion as a percentage of the lower bound. You can also specify an upper limit on the number of iterations allowed.

Limitations

Benders cannot model any of these features:

Automatic maintenance scheduling for planning alternatives

Multiple loading blocks for a unit

Must-run units

Spinning reserve

Probabilistic allocation of energy for storage units

Interconnections for reserve sharing

Interconnections for economy interchange

Detailed costs

Price elasticity

Fuel and emission constraints

6
DYNAMIC PROGRAMMING

Overview

The dynamic programming method is based on enumerating all possible planning alternative combinations in each year of the study period and selecting minimum cost transitions from year to year. In addition to finding the optimal expansion plan, the dynamic program (DP) can generate many suboptimal plans. The DP option is the most robust and detailed optimizing tool in EGEAS. Its major shortcoming is the large amount of computer time required to analyze a large number of planning alternatives.

Objective Function

For the DP, you have two choices for the objective function: present worth of revenue requirements and levelized system average rate. These quantities are described in the Objective Functions section of Chapter 4. The DP is capable of handling all the types of costs listed in Chapter 4, but the inclusion of some of the costs (e.g., detailed costs) depends on input options that you specify.

There are a few cases where EGEAS cannot guarantee an optimal solution. In general, these cases involve planning alternatives that have short operating lives or characteristics that change depending on the age of the unit. In addition, there are some limitations when the objective function is the levelized system average rate and there are DSM programs. For details on these cases, refer to the Limitations section later in this chapter.

In addition to the objective function used for optimization, an optional cost postprocessor evaluates expansion plans for economic feasibility. In particular, the postprocessor uses single-year and multiple-year average values of interest coverage ratio and earning assets ratio to determine which plans are feasible. For more details, refer to the Cost Postprocessor section.

Decision Variables

The decision variables for the DP are discrete (integer) variables representing the number of resources of each capacity planning alternative installed in each study period year. The DP can model thermal, limited energy, storage, and non-dispatchable technology alternatives.

The DP can also handle all the types of planning alternatives listed in the Decision Variables section of Chapter 4, including: installation of new resources, implementation of short-term (single-year) measures, unit replacements or refurbishments, phased installation of new resources, and unit retirements. With the DP, you can make two or more alternatives mutually-exclusive. Refer to the Dependency Constraints section for more details on the modeling of these alternatives.

Constraints

With the DP, EGEAS can impose the following types of constraints:

Reserve margin: Lower and upper limits on the reserve margin for a given study year. Both minimum and maximum constraints are applied together in every study year, or no reserve margin constraints are applied at all. The maximum reserve margin constraint can be useful in quickly eliminating states with too much capacity; however, if you want only the minimum constraint, you can set the upper bound very high.

Loss of load probability: Lower and upper limits on the LOLP for a given study year. You can specify both maximum and minimum constraints, maximum constraints only, or no constraints. As with the maximum reserve margin constraint, the minimum LOLP constraint can be used to eliminate states which have too much capacity, and hence are likely to be too expensive. The limits can be changed from year to year.

One other option is to constrain the relative reliability instead of the LOLP itself. Relative reliability gives the ratio of a benchmark value of LOLP to the LOLP for a particular year (see the Reliability Measures section of Chapter 2 for more details). If you choose the relative reliability option, the minimum value is always 1.0, meaning the LOLP should be no higher than the benchmark value.

Unserved energy: Upper limit on the expected unserved energy for a given study year. You can specify that the limit be applied to the annual unserved energy or to the unserved energy in each subperiod; the same limit (percentage of initial energy) applies in every subperiod of a year.

Tunnel: Lower and upper limits on the cumulative number of resources of a particular planning alternative that can be installed up through a given study year. The limits can vary from year to year and from alternative to alternative. An upper limit for one year is automatically applied to all previous years, unless smaller upper limits have been specified for those years. Similarly, a lower limit for one year automatically applied to all later years, unless larger lower limits have been specified for those years. Tunnel constraints are not allowed for single-year alternatives.

Alternative installation: Lower and upper limits on the number of resources of a particular planning alternative installed in a particular year. The limits can vary from year to year and from alternative to alternative; they are allowed for any type of alternative.

Fuel availability and minimum fuel usage constraints: essentially maximum and minimum limits on the amount of fuel, of particular types, used in a year.

Unit, site, and system emissions: upper limits on emissions, of particular types, in a year. The limits can be changed from year to year.

Emission rate reduction: Upper limit on the rate of emissions (e.g. tons CO_2 per GWH) produced by a particular group of units. Constraints can be satisfied by installing qualifying new "clean" generating units to the group of polluting units.

Renewable Portfolio Standard (RPS): Lower limit on the percent of generation provided by renewable technology (e.g. solar, wind) generating units.

After the DP generates the optimal and suboptimal expansion plans, it can evaluate these plans using a cost postprocessor (described in the Cost Postprocessor section). The postprocessor marks plans as feasible or infeasible. For an expansion plan to be marked feasible, it must satisfy the following constraints:

Earning assets ratio: Lower limit on the earning assets ratio for a given study year or group of years. The DP applies the constraints in every study year or not at all.

Interest coverage ratio: Lower limit on the interest coverage ratio for a given study year or group of years. The DP applies the constraints in every study year or not at all.

Percentage increase in system average rate: Upper limit on the percentage increase in system average rate over the previous year or over several years. The DP applies the constraints in every study year or not at all.

For more information on constraints, see the Constraints section in Chapter 4.

Methodology

The dynamic programming methodology involves the structuring of an optimization problem into multiple stages. These stages, which themselves are smaller optimization problems, are solved recursively to obtain a solution to the overall multi-stage problem. The optimization can start with either the first or last stage, using a forward or backward induction process, respectively. At each stage, a finite set of states, defined by one or more state variables, describes the possible choices. Each state contains the information necessary for making future decisions without regard to how the state was reached.

For the generation expansion optimization problem in EGEAS, each stage is a year of the study period. Each state is defined by the cumulative number of resources of each planning alternative installed in or before the current year. Figure 6-1 shows a set of states for a problem with two planning alternatives and a 3-year study period. In year 2, for example, there are 3 states: state (1,4) has 1 coal unit and 4 combined cycle oil units; state (1,5) has 1 coal unit and 5 combined cycle oil units; and state (0,7) has 7 combined cycle oil units. Each state indicates how many resources have been built, but not when they were built.

In Figure 6-1, the lines between the states indicate valid transitions. Because the numbers of alternatives are cumulative values, they cannot decrease from one year to the next. Thus, the transition from state (0,5) in year 1 to state (1,4) in year 2 is not valid. Each set of lines connecting a state in year 1 with a state in the last study year represents an expansion plan. For example, the lines connecting states (0,4), (1,5), and (1,7) represent a plan in which 4 combined cycle oil units are installed in the first year, a coal unit and one more combined cycle oil unit are installed in the second year, and 2 more combined cycle oil units are installed in the third year.

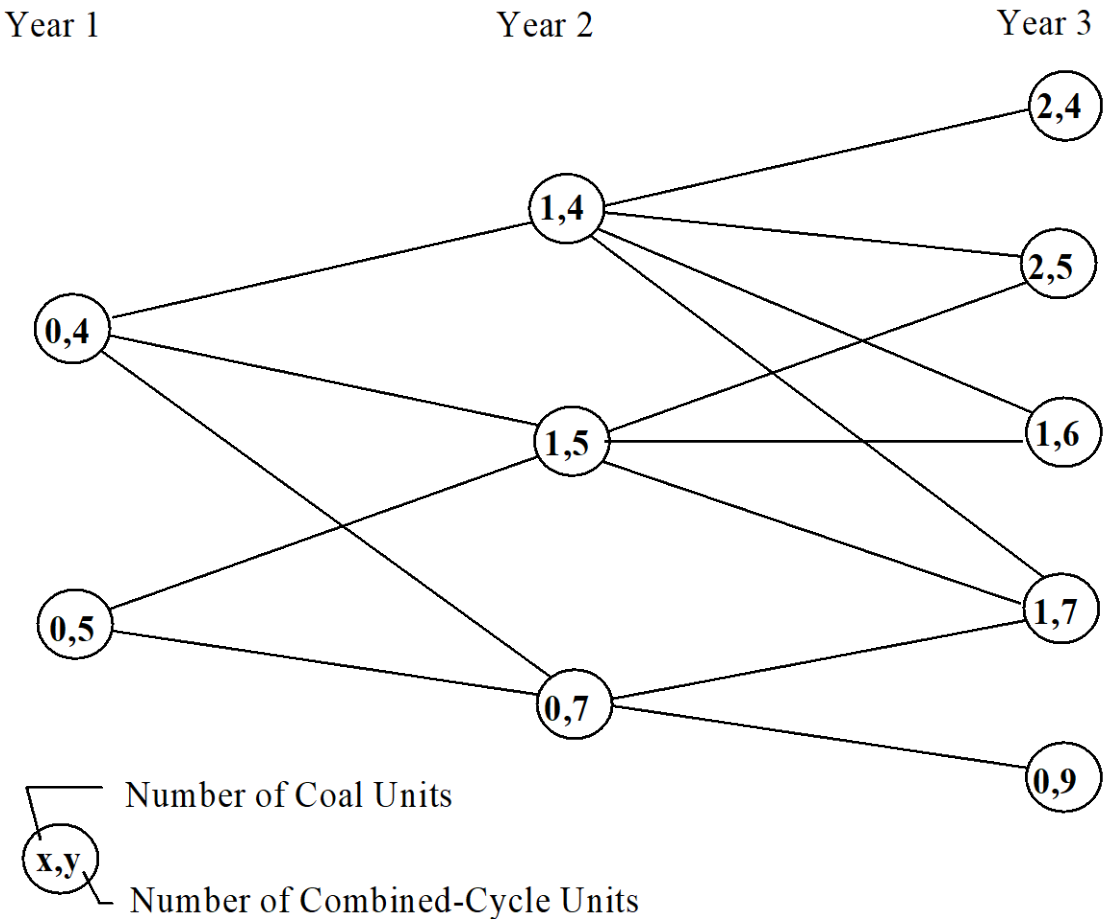

Figure 6-1
Dynamic Program States

Note that single-year alternatives, those that can operate for only one year, are treated differently from other alternatives. For single-year alternatives, the state definition includes the number of units installed in that year rather than the cumulative number of units. Thus, in Figure 6-1, the transition from state (0,5) in year 1 to state (1,4) in year 2 would be valid if the second alternative was a single-year alternative. In that case, the DP would install 5 units of the second alternative in year 1 to operate only in year 1 and then would install 4 units in year 2 to operate only in year 2.

The basic methodology for the DP option involves five main steps. The first step is forward dynamic programming (forward DP), which implements a forward induction process. In this step, the DP generates and evaluates all the states in each study year, starting with the first year. For each state, the DP determines an optimal path for reaching the state from a state in the previous year. The second step is backward dynamic programming (backward DP), which implements a backward induction process. This step, which is optional, is very similar to the forward DP step, but starts with the last study year. The backward DP step is computationally much faster than the forward DP step because it can use many of the results calculated by the forward DP. While the DP can determine the optimal expansion plan without this step, the backward DP lets the DP generate a better set of suboptimal plans. The third step is construction of the optimal expansion plans. In this step, the DP uses the information gathered in the first two

steps to determine the optimal and suboptimal expansion plans. The fourth step is the cost postprocessor. In this optional step, each of the expansion plans generated in the previous step(s) are evaluated to determine whether they satisfy the economic feasibility constraints. You specify whether infeasible plans should be retained or discarded. The fifth step simulates the expansion plan(s) retained during the optimization step and generates information for reporting.

State Generation

During forward DP processing, the DP must generate a list of possible states for each year. In the absence of any constraints, the DP would try all possible combinations of the planning alternatives specified. The DP arbitrarily places an upper limit of 99 on the number of units for any alternative to keep the number of states finite. Even with this limit, however, the number of states can become very large. Four mechanisms control the state generation process, helping to limit the number of states: tunnel constraints, upper reserve margin constraints, limits on the number of superfluous units, and dependency constraints.

First, the tunnel constraints specified place lower and upper limits on the numbers of units for each planning alternative. For example, if there are two planning alternatives and you limit the first alternative to a maximum of 2 units and limit the second alternative to a minimum of 5 and a maximum of 6 units, then the DP generates the following states:

(0,5) (1,5) (2,5) (0,6) (1,6) (2,6)

where the numbers in parentheses are the numbers of units of the first and second alternative, respectively. When generating states, the DP starts with all alternatives at their minimum values and then increases the number first for alternative 1, then for alternative 2, then for alternative 3, etc. The order in which you select the alternatives affects the order in which the states are generated, but does not affect the optimization unless there are more states than the DP can handle (see **State Compression** section).

Second, the upper reserve margin constraint keeps the DP from generating states with excess numbers of resources. If the DP finds that the reserve margin is too large for one state, then it will not generate states with even more resources. Thus, in the example given above, the DP would not generate state (2,6) if the reserve margin for state (1,6) exceeded the upper reserve margin limit. However, the DP would generate state (0,6) even if state (2,5) had too high a reserve margin because the reduction of 2 units of the first alternative might bring the reserve margin down below the upper limit with enough room for the additional unit of the second alternative.

Third, the DP rejects states with superfluous units. A superfluous, or extra, unit is one that can be removed from a state without affecting its feasibility. In the example above, if state (0,5) is feasible, then 1 unit (of alternative 1) can be removed from state (1,5); thus state (1,5) has one superfluous unit. Similarly, state (2,6) has 3 superfluous units because 2 units of alternative 1 and 1 unit of alternative 2 can be removed. You can impose an upper limit on the number of superfluous units allowed. If the limit were 2 in the above example, then state (2,6) would be eliminated. If the limit were 1, then states (2,5), (1,6), and (2,6) would all be eliminated.

In most cases, economics dictate that a utility build only enough resources to satisfy the reliability constraints, so eliminating states with superfluous units does not affect the optimization. However, the optimal solution might contain extra units in some years for one of the following reasons:

If you have a planning alternative that installs a relatively small resource (e.g., a DSM program), that alternative might not be enough by itself to displace another alternative, and yet it might result in a lower cost expansion plan.

If you have a planning alternative that retires an existing unit, it is probably superfluous, because it provides no new energy; yet, it might be very cost effective.

If you are modeling the phased installation of a unit, the alternatives representing the early stages may be superfluous; yet they are needed to allow the DP to proceed to the productive and cost-effective later phases.

If a resource has a high immature forced outage rate, it might be advantageous to install the resource a year or two before it is needed so it can operate more reliably when needed.

If capital costs are escalating more rapidly than other costs, it might be cost effective to install a unit in an earlier year when the capital costs are lower.

If a new resource is much more efficient or uses a less expensive fuel than most existing resources, then the operational savings of displacing old inefficient resources might outweigh the capital costs of installing new capacity.

If you have restricted a resource to installation only in certain years, it may be cost effective to install the resource when it is available for installation even though it is not needed until a later year.

Thus, you should choose the maximum number of superfluous units carefully. If it is too small, the optimum solution may be lost; if it is too high, then the DP may evaluate many states that are far too expensive for inclusion in an optimal expansion plan.

Usually, you should consider allowing a superfluous unit for each alternative that retires an existing unit and each new unit that is small in size compared to other alternatives (e.g., DSM programs). Note that you can specify separate limits on the number of superfluous units for each alternative type.

Fourth, the generation of states can be affected by dependency constraints. These constraints provide linkages between two or more types of planning alternatives in several different ways, such as:

If you have specified that two or more alternatives are mutually-exclusive, no state will be generated that includes more than one from that group of alternatives.

If you have specified that one alternative is a prerequisite for another, the DP will not generate a state that includes an alternative without its prerequisite.

If you have specified that one alternative must be installed with or soon after another alternative, the DP will not generate a state that does not meet that requirement.

For more details, see the Dependency Constraints section later in this chapter.

For a particular analysis, the number of states generated can easily reach into the thousands. For example, with just 10 alternatives limited to 0 or 1 units each, there are 1,024 possible states; with 20 alternatives limited to 0, 1, 2, or 3 units each, there are over 1 trillion possible states. The DP, however, limits the number of feasible states (states that satisfy all constraints). The actual limit depends on your data and the modeling options you have selected; more details are given in the EGEAS User's Guide. If this limit is reached, the DP must stop generating states, raising the possibility of missing the optimum solution. Thus, you should be very careful in choosing

reasonable reliability constraints, tunnel constraints, and limits on superfluous units. If your limits are too loose, the number of states becomes very large, the running time can extend to days, and you may lose the optimum solution through having too many states. If your limits are too strict, you may exclude some low-cost states and hence, lose the optimum solution. Refer to the practical guidelines section for more details.

State Evaluation

During forward DP processing, the DP evaluates every state generated. The two objectives of state evaluation are to eliminate infeasible states (i.e., states that do not satisfy all the constraints) and to calculate costs used for determining the optimal expansion plan. There are five main steps in evaluating a state:

Superfluous unit check: The DP compares the state with all previously-generated feasible states for the current year, if any. If the set of resources for any of these states is a subset of the resource for the current state, then the current state has superfluous units. The number of such extra units may be different for comparisons with different states. For example, state (2,6) has 3 extra units when compared with (0,5) but only 1 extra unit when compared with (1,6). If the number of superfluous units in any comparison exceeds the upper limit, the current state is infeasible. The DP checks upper limits on the number of superfluous units of particular alternatives, if any, as well as the overall upper limit.

Reserve margin constraints: The DP calculates the reserve margin and checks whether it is within the acceptable range of values. If the reserve margin is too low, the state is infeasible; if it is too high, the state is infeasible and the DP will not generate further states with more units added.

Loss of load probability constraint: The DP calculates the loss of load probability (LOLP) and checks whether it is within the acceptable range of values. If not, the state is infeasible. If you have specified benchmark data and requested a relative constraint, then the DP calculates relative reliability; the state is infeasible if the relative reliability is less than 1.0. Note that the relative reliability constraint replaces the LOLP constraint.

Pathway constraint: The DP searches through all the feasible states for the previous year, looking for states that provide valid transitions to the current state. Because the numbers of resources are cumulative, a transition is valid only if the resources comprising the older state are a subset of the resources comprising the newer state. No checking is needed for a single-year alternative, however, because the number of units in one year is independent of the number of units in any other year. If there are no valid transitions to the current state, the state is infeasible because it cannot be reached from the previous year. For example, if every state in year 2 has at least 4 units of the second alternative, then a state in year 3 with only 3 units of the second alternative is infeasible. Note that this step is bypassed in the first study period year.

Unserved energy constraint: The DP determines the production cost and the unserved energy. If the unserved energy exceeds the upper limit, the state is infeasible. If subperiods are modeled, the DP calculates the cost and unserved energy for each subperiod and sums these subperiod values to get the corresponding annual totals. With subperiod modeling, you can specify that the unserved energy constraint be applied in each subperiod or on the annual total. With subperiod constraints, the state is infeasible if the unserved energy exceeds the upper limit in at least one subperiod. For feasible states, the annual production cost total is saved for use in determining the optimal expansion plan.

Emission constraint: If you have specified a system-wide constraint on an emission type, you have the further option of specifying that states which cannot meet the constraint be rejected. In this case, the DP checks the emissions after production costing is complete; if the emissions are too high, the state is rejected as infeasible.

State Compression

EGEAS does not limit the number of states that may be generated in a year, but does limit the number of feasible states. The DP must retain information about each feasible state, so computer storage requirements dictate how many feasible states are allowed. Actually, the DP imposes two limits: a limit on the number of feasible states generated in the current year and a limit on the number of feasible states retained for determining optimal plans. If the second limit is less than the first, then the DP uses state compression when necessary to select which states should be retained.

During forward DP processing, the DP keeps track of the number of feasible states found. If the DP reaches the upper limit for feasible states generated in a year, then it stops generating states for that year. In this situation, optimality may be lost because the best expansion plan might include a state that was not yet generated. Before the DP continues processing, it prints a warning message to alert you of this condition.

If the number of feasible states actually generated in a year exceeds the limit on retained states, then the DP uses state compression to eliminate some of the feasible states. EGEAS allows you to specify several criteria for determining which states should be eliminated. There are two phases, each with different options. If the number of states is reduced sufficiently during the first phase, then the DP does not use the second phase. For the first phase of state compression, the options are:

Take no action: This option allows the DP to proceed immediately to the second phase of state compression.

Reduce superfluous units: This option eliminates the states with the most superfluous units. States with the same number of extra units are eliminated as a group. For example, if the maximum number of superfluous units is 4, then the DP first eliminates all states with 4 extra units, then all states with 3 extra units, and so forth until the number of states remaining is within the allowed limit. Because states are eliminated in groups, the remaining number of states may be less than the allowed limit. Note that this option never eliminates states with no extra units.

Eliminate superfluous units: This option eliminates every state with at least one superfluous unit. As with the previous option, the remaining number of states may be less than the allowed limit. This option has the added effect of reducing the upper limit on superfluous units to zero for all later study years. Thus, in all later years, any state with an extra unit will be infeasible. For the second phase of state compression, the options are:

Eliminate units arbitrarily: This option retains the first feasible states generated, up to the number of states that can be retained, and eliminates the remainder. Note that the effect of this option depends on the order in which you specify the planning alternatives.

Reserve margin: This option sorts states by reserve margin and retains those with the smallest values.

Alternative units: For this option, you select one of the planning alternatives. The DP sorts states by the cumulative number of units of this alternative and retains the states with the smallest number of resources.

Committed cost: This option sorts states by their committed cost (described in the **Forward Dynamic Programming** section) and retains those with the lowest cost.

While some state compression options may yield better results than others under certain conditions, none of the options is guaranteed to retain the states needed to generate the optimal expansion plan. Thus, when the DP uses state compression, it may not find the optimal expansion plan. For more information on conditions for optimality, see the **Practical Guidelines** section.

Forward Dynamic Programming

The forward DP steps through the years of the study period in order, starting with the first year. For each year, the forward DP generates all possible states, evaluates them, and, if necessary, compresses them using the methods described in the preceding sections. For each feasible state, the forward DP also determines a committed cost and a backward pointer.

The committed cost for a state is the cost of the least expensive path leading to that state. More precisely, the committed cost includes components of the objective function related to costs that have been calculated up through the current state. For present worth of revenue requirements, the committed cost includes various cost components, discounted to the base year. For levelized system average rate, the committed cost includes various cost components divided by system sales and then discounted to the base year (levelization is applied later). For either objective function, the following components are included:

Operating costs: Includes fuel, fixed O&M, and variable O&M costs for all resources: existing, committed, and generic. Also includes detailed costs that are modeled as fixed or variable O&M, purchased energy costs (less sales revenues), emission allowance costs (less credits), and unserved energy costs.

Capital costs: Includes the capital fixed charges for installing new resources and capital expenditures for resources already installed. Also includes detailed costs that are modeled as capital costs.

Costs for "non-EGEAS" assets: If the objective function is levelized system average rate, the costs associated with "non-EGEAS" assets are included. Fixed capital charges and operating costs are added and other electric revenues are subtracted.

Most costs are summed up over all the years from the beginning of the study period up to and including the current year for the state. For the last study period year, costs for the extension period are included as well. However, there are a few cases where the costs are projected beyond the current year, all the way to the end of the extension period, even for states in the early years:

Capital fixed charges, and detailed costs modeled as capital expenses, for planning alternatives only.

Detailed costs modeled as capital expenses for existing or committed units, but only if the objective function is revenue requirements.

Detailed costs modeled as fixed O&M with an expenditure pattern, but only for planning alternatives that cannot be forced into retirement by a later alternative.

If the objective function is revenue requirements, the committed cost for a state in the last study year is the cost of the least expensive plan terminating at that state. If the objective function is levelized rate, the committed cost for an end-year state, when levelized, yields the levelized system average rate for the least cost plan that terminates at that state. Note that the levelization process simply involves multiplying the committed cost by a constant. thus, the plan with the lowest committed cost is also the plan with the lowest levelized system average rate.

There are several reasons why some components of the committed cost are summed over all years and others are summed only through the current study year:

Production costs cannot be determined for future years until the production costing dispatch is performed in those later years; hence, operating costs for later years are usually not included.

In most cases, capital costs for later years are predetermined when a resource is installed; hence, capital costs for later years can be included if needed.

When the objective function is levelized system average rate, components of the objective function for later years depend on the system sales, which in turn depend on the level of DSM output. Because the DSM contributions for later years is usually not known until the DP gets to those years, the later year components are excluded if possible.

In some cases (e.g., capital costs for existing and committed units), the results would be the same either way and the costs are included or excluded according to convenience.

To calculate the committed cost for a state, the DP needs to know the path taken to reach the state. Each different path includes a distinct set of previous year states, so the production costs for earlier years vary with the path chosen. Each path also determines a distinct pattern of installation dates for the planning alternatives, so the capital costs for the generic resources also vary with the path chosen. To determine the best path for a particular state, the DP calculates the committed cost for different paths and chooses the path with the lowest cost.

The backward pointer for a state indicates the state in the previous year that is included in the best path to the current state. The complete path to a particular state can be found by following the chain of backward pointers all the way back to a state in the first study year (there are no backward pointers for states in the first year). When processing a state, the DP needs to look back only one year because the DP has already determined the best paths between states in the earlier years.

As noted earlier, the committed cost for a state in the last year is the total cost of the best expansion plan terminating in that state. To determine the optimal expansion plan, the forward DP simply follows the backward pointers from the state in the last year with the smallest committed cost. By following the pointers from other states, the DP can generate additional suboptimal plans. Although the suboptimal plans tend to be good plans, they are not guaranteed to include the second best plan, because the two best plans might end at the same state in the last year.

Figure 6-2 demonstrates how the DP calculates committed costs and determines backward pointers. In this example, the objective function is present worth of revenue requirements. For each state, the figure shows the committed cost (CC) and the components of that cost:

PC = Production cost for the current year (for the last study year, the extension period production cost is included)

FIX = Fixed costs (capital fixed charges) for units installed in the current year (as determined by the backward pointer)

CCX = Committed cost for the previous year state indicated by the backward pointer (includes production costs for earlier years and fixed costs for units installed before the current year)

To illustrate the calculations, consider state (0,4) in year 1. The production cost is 586 (all costs are in millions of discounted dollars). The fixed cost for the four combined cycle oil units, projected through the end of the extension period, is 445. Because this is the first year, there is no backward pointer or previous year committed cost, so the committed cost for state (0,4) is 586 + 445 = 1031.

For state (1,5) in year 2, the production cost is 536. The DP has two choices for backward pointers from this state: to (0,4) or to (0,5). For a path that included state (0,4), the committed cost would include the production cost (536), the fixed cost of 1 coal unit and 1 additional combined cycle oil unit (1126) and the committed cost for state (0,4) (1031); this totals to 2693. For a path that included state (0,5), the committed cost would include the production cost (536), the fixed cost of 1 coal unit (1023), and the committed cost for state (0,5) (1137); this totals to 2696. Because the committed cost for the path through state (0,4) is smaller, the backward pointer from state (1,5) points to state (0,4).

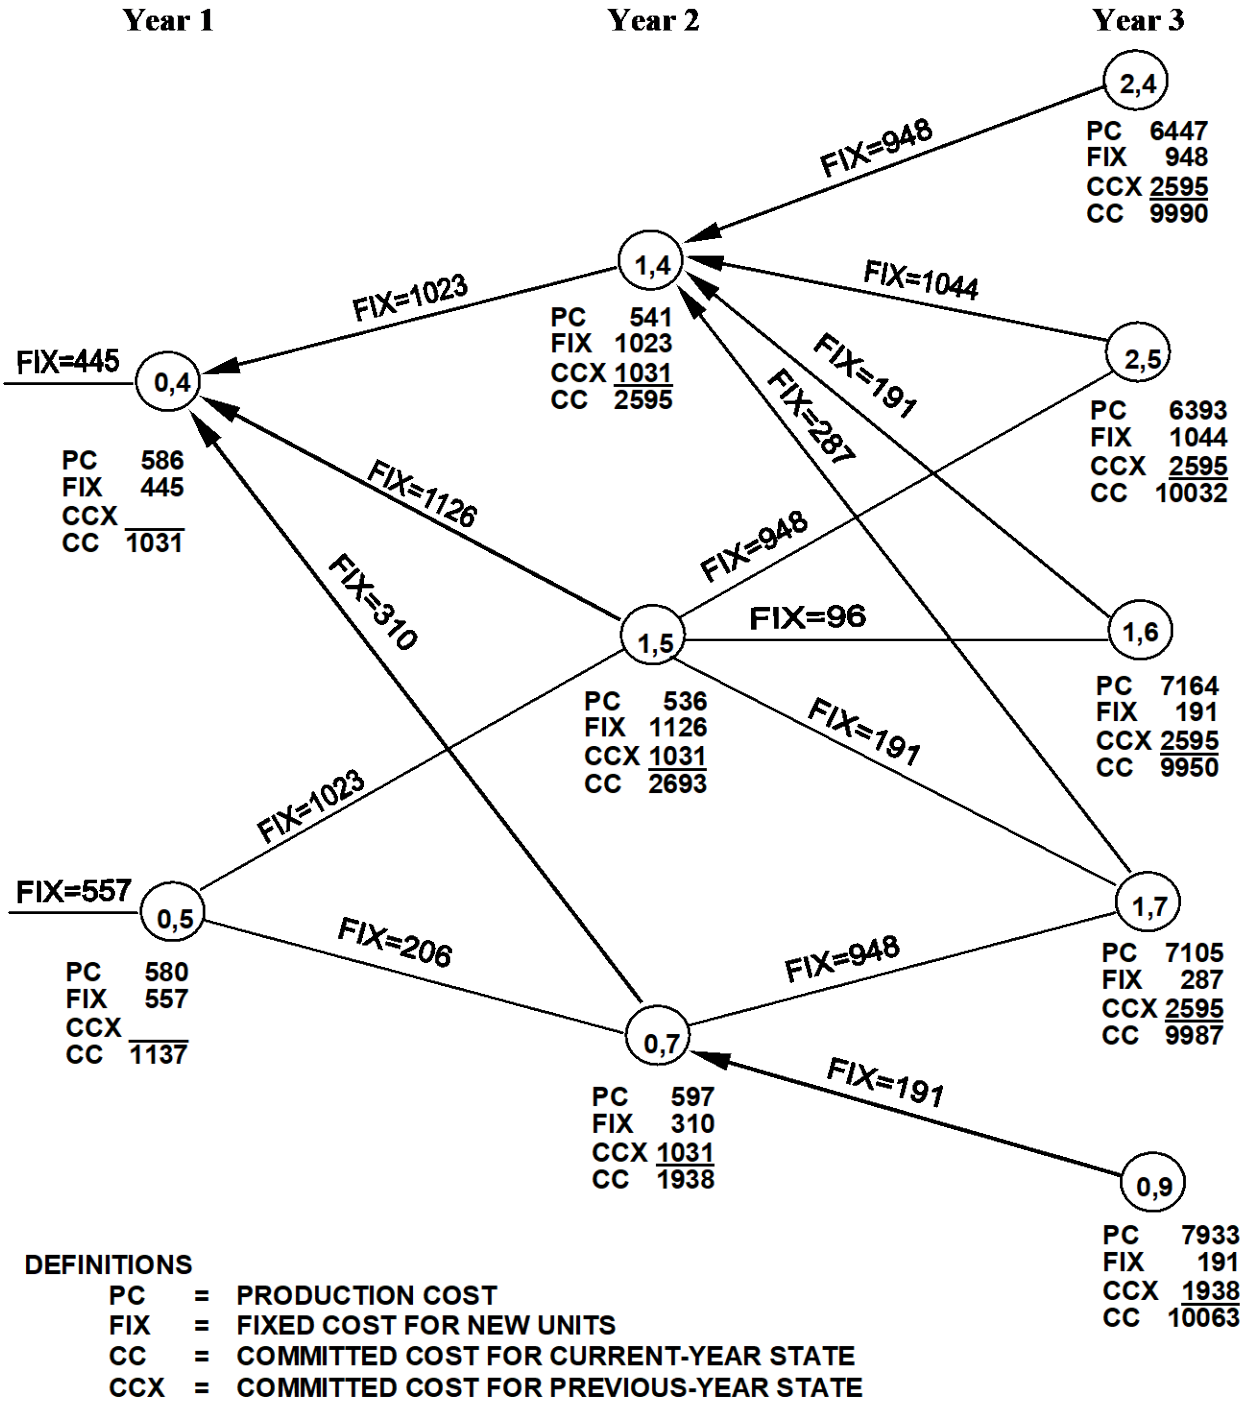

Figure 6-2
Committed Costs and Backward Pointers

The committed cost for a state in year 3 represents the total cost of the best plan that ends in that state. Note that production costs are much higher in year 3 because they include costs for a multi-year extension period. State (1,6) has the lowest committed cost (9950), so the optimal expansion plan is determined by following the backward pointers to state (1,4) in year 2 and state (0,4) in year 1. Other (suboptimal) plans are represented by the other states in year 3. Table 6-1 shows all the possible expansion plans for this case and indicates which plans are found by the forward DP. As expected, the forward DP finds the best plan. However, it misses the second and third best plans because they both end at the same state (1,6) as the best plan. For larger, more realistic cases, the forward DP usually finds many plans that cost nearly the same as the optimal plan.

Table 6-1
Expansion Plans from Forward DP

Year 1	Year 2	Year 3	Total Plan Cost	Ranking	Forward DP
(0,4)	(1,4)	*(1,6)	9,950	1	X
(0,4)	(1,5)	(1,6)	9,953	2	
(0,5)	(1,5)	(1,6)	9,955	3	
(0,4)	(1,4)	*(1,7)	9,987	4	X
(0,4)	(1,5)	(1,7)	9,989	5	
(0,4)	(1,4)	*(2,4)	9,990	6	X
(0,4)	(0,7)	(1,7)	9,991	7	
(0,5)	(1,5)	(1,7)	9,992	8	
(0,5)	(0,7)	(1,7)	9,994	9	
(0,4)	(1,4)	*(2,5)	10,032	10	X
(0,4)	(1,5)	(2,5)	10,034	11	
(0,5)	(1,5)	(2,5)	10,037	12	
(0,4)	(0,7)	*(0,9)	10,063	13	X
(0,5)	(0,7)	(0,9)	10,065	14	

* = optimal forward DP plan through this state

Backward Dynamic Programming

The backward DP is very similar to the forward DP except that it starts with the last study year and works back to the first year. The backward DP operates much faster, however, because it can use many of the results from the forward DP. In particular, the backward DP does not need to determine which states are feasible or calculate production cost and reliability. While the forward DP calculates the committed cost and determines a backward pointer for a state, the backward DP calculates a cost-to-go and determines a forward pointer. The cost-to-go includes the same component costs as the committed cost used with forward dynamic programming, with the following differences:

Operating costs: These are summed over all the years from the current year to the last year of the study period; the extension period costs are always included.

Capital costs: For existing and committed resources, these costs are summed over the years from the current year to the end of the extension period. For planning alternatives, the costs are summed over all years, for alternatives that are installed in later years (after the current year). Capital costs for planning alternative units included in the current state are NOT included, because the installation dates are not known until the DP works back to an earlier year.

Costs for "non-EGEAS" assets: These costs are summed over all years from the current year to the end of the extension period.

The backward DP calculates the cost-to-go for all possible transitions to states in the following year and chooses the least expensive option. A forward pointer indicates which state in the following year was chosen. Note that there are no forward pointers for the last study year and that costs for states in the last year contain no fixed costs.

For every study year except the last, the backward DP determines the best expansion plan that passes through each state. The last year is not included because the forward DP has already determined the best plan through each state in that year. The best plan for a particular state is the one determined by following the backward pointers (from the forward DP) to the first study year and following the forward pointers (from the backward DP) to the last study year. The total cost for such a plan is the sum of the committed cost and the cost-to-go, less those cost components for the current year which are included in both the committed cost and the cost-to-go (to avoid duplication). The current-year costs which are subtracted include:

Production costs, including detailed costs modeled as variable O&M

Capital fixed charges for existing and committed units

Detailed costs modeled as fixed O&M, except for planning alternatives where there is an expenditure pattern and the alternative cannot be forced into retirement

Detailed costs modeled as capital costs, but only for existing or committed units and only when the objective function is the levelized system average rate

Costs for "non-EGEAS" assets, when the objective function is levelized system average rate

You may choose whether the backward DP should be performed. The backward DP never changes the optimal expansion plan, so it is useful only for determining a wider range of suboptimal plans. The forward DP is guaranteed to produce only the best plan; it may not find the second best plan. The forward and backward DP together are guaranteed to produce the two best plans; they may not find the third best plan.

Figure 6-3 demonstrates how the DP calculates the cost-to-go and determines the forward pointer for a state. In this example, the objective function is present worth of revenue requirements. For each state, the figure shows the cost-to-go (CG) and the components of that cost:

PC = Production cost for the current year (for the last study year, the extension period production cost is included)

FIX = Fixed costs (capital fixed charges) for units installed in the year after the current year

CGX = Cost-to-go for the next year state indicated by the forward pointer (includes production costs for later years and fixed costs for units installed after the current year)

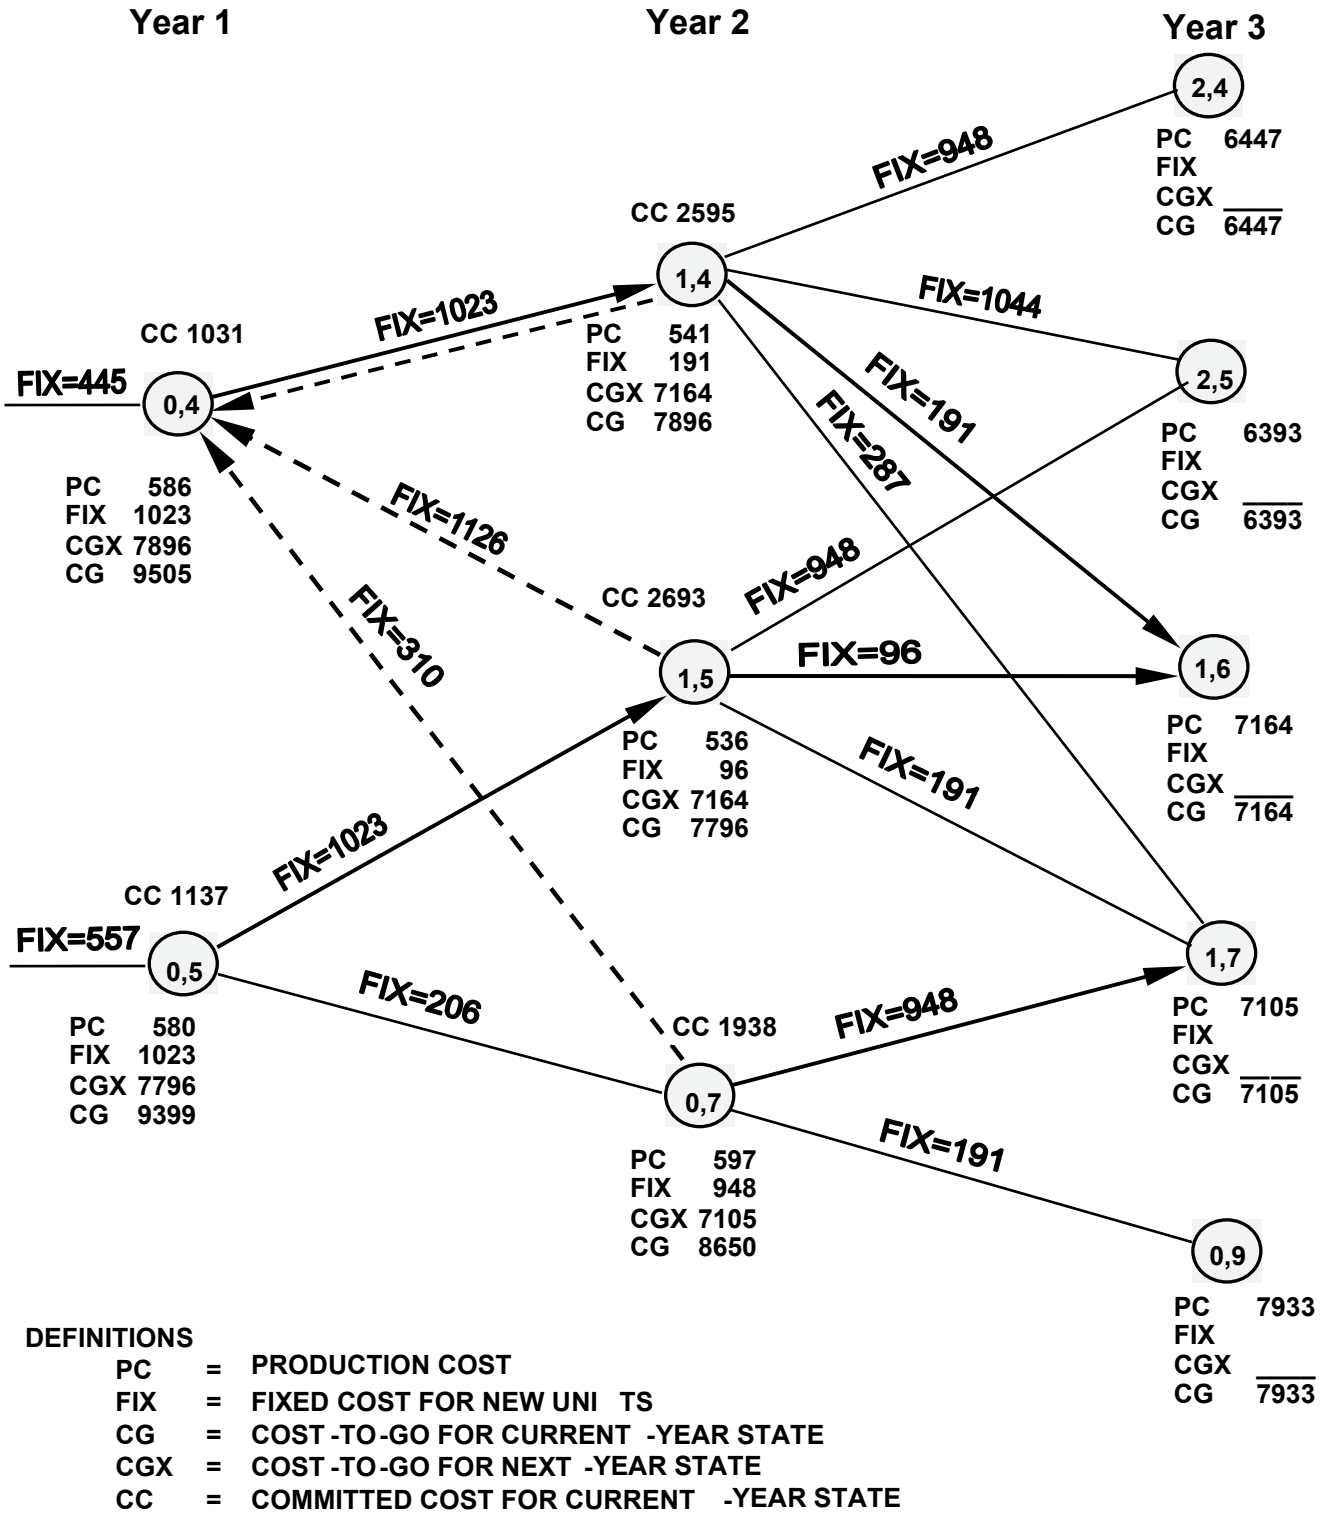

Figure 6-3
Costs to Go and Forward Pointers

To illustrate the calculations, consider state (1,7) in year 3. The cost-to-go is simply the production cost summed over year 3 and the extension period (7105). The best plan through this state was determined by the forward DP because this is the last study year.

For state (0,7) in year 2, the production cost is 597. The DP has two choices for forward pointers from this state: to (1,7) or to (0,9). For a path that includes state (1,7), the cost-to-go would include the production cost (597), the fixed cost of 1 coal unit added in year 3 (948), and the cost-to-go for state (1,7) (7105); this totals to 8650. For a path that includes state (0,9), the cost-to-go would include the production cost (597), the fixed cost for 2 additional oil units (191), and the cost-to-go for state (0,9) (7933); this totals to 8721. Because the cost-to-go for the path to state (1,7) is smaller, the forward pointer from state (0,7) points to state (1,7).

The best plan through state (0,7) includes state (0,4) in year 1 (indicated by the backward pointer) and state (1,7) in year 3 (indicated by the forward pointer). The total cost of this plan is the sum of the committed cost (CC = 1938) and the cost-to-go (CG = 8650) less the production cost (PC = 597); this cost is 9991. Table 6-2 shows the expansion plans found by the backward DP.

Table 6-2
Expansion Plans from Forward and Backward DP

Year 1	Year 2	Year 3	Total Plan Cost	Ranking	Forward DP	Backward DP
+(0,4)	+(1,4)	*(1,6)	9,950	1	X	X
(0,4)	+(1,5)	(1,6)	9,953	2		X
+(0,5)	(1,5)	(1,6)	9,955	3		X
(0,4)	(1,4)	*(1,7)	9,987	4	X	
(0,4)	(1,5)	(1,7)	9,989	5		
(0,4)	(1,4)	*(2,4)	9,990	6	X	
(0,4)	+(0,7)	(1,7)	9,991	7		X
(0,5)	(1,5)	(1,7)	9,992	8		
(0,5)	(0,7)	(1,7)	9,994	9		
(0,4)	(1,4)	*(2,5)	10,032	10	X	
(0,4)	(1,5)	(2,5)	10,034	11		
(0,5)	(1,5)	(2,5)	10,037	12		
(0,4)	(0,7)	*(0,9)	10,063	13	X	
(0,5)	(0,7)	(0,9)	10,065	14		

* = optimal forward DP plan through this state

+ = optimal backward DP plan through this state

Construction of Expansion Plans

The forward DP generates committed costs and backward pointers, which determine the least expensive plans ending in each state of the last year. The backward DP generates costs-to-go and forward pointers, which with the committed costs and backward pointers, determine the least

expensive plans through states in years other than the last year. The DP sorts these plans in order of increasing total cost and selects the best plans for detailed simulation and reporting. You specify the number of plans to be generated.

The backward DP typically generates many duplicate plans because the same plan may be the best plan through several states. In particular, the optimal expansion plan must be the best plan for each of the states in the path for that plan. For example, in Table 6-2, plan number 1 is the best plan for state (0,4) in year 1, for state (1,4) in year 2, and for state (1,6) in year 3.

Simulation of Expansion Plans

After optimization is complete and the cost postprocessor has evaluated the plans, the DP simulates the optimal and suboptimal plans for reporting purposes. Simulation includes calculating reliability measures and capital, operating, and total costs. While these quantities were calculated previously during optimization, the detailed results (e.g., individual unit generation, subperiod unserved energy) for all the infeasible and feasible states contain too much information to be saved efficiently. The DP simulates the number of plans specified. However, the number may be lower if the DP is unable to find the requested number of unique plans or if the number is reduced by the cost postprocessor.

Installation Year Estimation

During the forward DP portion of the optimization, the DP calculates the production cost for each of the generated states. At this point in the optimization, the DP has not yet constructed any plans, so each state is an independent entity. The DP knows which planning alternatives are included in any given state, but does not know when the resources were installed. Indeed, the same state may become a part of several different plans, with different sets of installation dates. Thus, the DP cannot model accurately any parameter whose value depends on the age of a planning alternative. Examples of such parameters include the capacity of a resource, the forced outage rate of a resource (which may be higher during the early years of a unit's operation), and maintenance requirements (which may be high in years where major components are overhauled). The most serious effect is on unit retirement. While most new resources have operating lives which would extend beyond the end of the study period, a planning alternative can represent a purchase contract or some other capacity option with a short duration. If the DP does not know when a resource with a fixed operating life was installed, it does not know whether it has yet retired.

There is no satisfactory solution to this problem. The limitation is inherent in the formulation of the problem and the dynamic programming methodology. However, the DP must use some installation date for each generic resource so it can calculate the time-dependent characteristics. The procedure for choosing installation dates is as follows:

1. For the first study year, assume that all planning alternatives are installed in that year.
2. After state evaluation is complete for a particular year, determine the number of resources for an "average" state. For a particular planning alternative, the number of installed resources is averaged over all the feasible states in the current year. This average value is rounded to the nearest integer value and assigned to the "average" state. (In the average, each state is weighted by the inverse of its committed cost, so less expensive states have a higher weight than more expensive states.)

3. For a state in the second or any later year, determine installation dates assuming a path through the "average" state in each of the previous years. If there are more resources in an "average" state than in the current year state, take the earliest installation years.

As an example of installation year estimation, consider again the problem illustrated in Figures 6-1 and 6-2. For the first year, the DP determines the "average" state, based on state (0,4) with a committed cost of 1031 and state (0,5) with a committed cost of 1137, as follows:

Year 1, alternative 1: both states have 0 units

Year 1, alternative 2: $\frac{4/1031+5/1137}{1/1031+1/1137}=4.48$

Thus, the "average" state for year 1 is (0,4). Similarly, for year 2 the weighted averages over states (1,4), (1,5), and (0,7) yield an "average" state of (1,6). Realizing that these averages reflect cumulative numbers of units, the DP assumes the following numbers of units installed in each year:

	Year 1	Year 2
Alternative 1 (coal)	0	1
Alternative 2 (oil)	4	2

For state (2,5) in the last year, the DP assumes that 4 combined cycle oil units are installed in year 1 and that 1 coal unit and 1 combined cycle oil unit are installed in year 2. The remaining (second) coal unit is installed in year 3; the sixth combined cycle oil unit in the "average" state for year 2 is not needed.

The DP uses the installation year estimation procedure given above during the optimization. If resource characteristics change significantly, the optimization may be based on incorrect costs and the correct optimal solution may not be found. During simulation of the final plans, however, the installation dates are known and the plans are simulated correctly. Thus, the cost determined during final simulation may be different from that used during optimization and the optimal and suboptimal plans may be out of order according to their correct costs.

Because the DP cannot properly model changing resource characteristics for planning alternatives, you should avoid changing resource characteristics according to the age of the resource. Instead, you should set each characteristic at its average value over the life of the resource. In addition, you should not try to model units with short operating lives (lives shorter than the length of the study period). Note, however, that the DP does handle seasonal changes properly.

Installation year estimation is not a problem for committed resources because their installation dates are fixed. It is also not a problem for a planning alternative that is restricted to installation in only a single year.

Suboptimal Plans

One of the advantages of the DP is its ability to generate suboptimal expansion plans. In some cases, the costs of several plans may be sufficiently close that a utility chooses to follow a suboptimal plan rather than the optimal plan because of factors not considered by EGEAS. However, many of the best plans generated by the DP are often identical in the early study years,

with small variations in the later years where cost differences are less important because of cost discounting. For such cases, the DP provides you with little choice in which resource should be installed first. To overcome this difficulty, the DP can eliminate plans that are trivial variations of some better plan. The DP treats several plans as trivial variations of one another if they are identical in all but the last N years, where you specify the value of N. When the DP finds such a group of plans, it keeps the plan with the lowest cost and discards the others. The trivial variation feature is optional.

Dependency Constraints

In the DP, dependency constraints are used to model a variety of options for generation expansion planning. Among the capabilities available with dependency constraints are:

optimize retirement of an existing resource

optimize replacement of an existing resource with a new resource

make installation of one planning alternative contingent upon the installation of another (prerequisite) alternative

make two or more planning alternatives mutually-exclusive

optimize installation of resources that can operate for only a single year

allow planning alternative resources to be installed in groups, with a different capital cost for the first resource of each group

Prerequisite Alternatives

In each year of the study period, the dynamic program (DP) generates all possible combinations of planning alternative, subject to lower and upper tunnel constraints. The number of resources for a planning alternative is the cumulative number of resources added through the current year. Note that planning alternatives that retire before the end of the study period are not subtracted from this count. If one planning alternative requires another alternative as a prerequisite, the DP checks whether the required prerequisites have been installed. Any state that does not have the required prerequisites is discarded immediately -- it is not included in the number of states generated.

You can specify that N units of a prerequisite alternative are required for each unit of the follow-on resource, or that a single prerequisite alternative is sufficient to allow installation of up to N follow-on resources, where N can range from 1 to 99. Thus, for example, you could specify that 2 combustion turbines (prerequisite alternative) be installed before a combined cycle unit (follow-on alternative) is installed in a phased-installation arrangement. Or, you could specify that a DSM program with high marketing and startup costs (prerequisite alternative) be installed first, followed by up to 5 units of another DSM programs (follow-on alternative).

In addition, you may specify up to four different prerequisite planning alternatives for a given alternative. In this case, EGEAS will install a resource of the current alternative only when all of the prerequisite conditions are met.

As an example, consider a problem with two planning alternatives. Assume that the maximum number of resources is 4 for each type. Assume further that 2 units of alternative 1 are required as prerequisites for each unit of alternative 2. The DP would generate the following states:

1	(0,0)	(0,1)*	(0,2)*	(0,3)*	(0,4)*

2	(1,0)		(1,1)*		(1,2)*	(1,3)*	(1,4)*
3	(2,0)	6	(2,1)		(2,2)*	(2,3)*	(2,4)*
4	(3,0)	7	(3,1)		(3,2)*	(3,3)*	(3,4)*
5	(4,0)	8	(4,1)	9	(4,2)	(4,3)*	(4,4)*

The states marked with an asterisk (*) would be discarded because they do not have the necessary prerequisites. The DP would report that 9 states were generated.

Associated with prerequisite alternatives are minimum and maximum lag year options. The minimum lag year option specifies the minimum number of years between the installation of the last prerequisite unit and the installation of the follow-on alternative. The options are:

Option 0: The follow-on alternative may be installed in the same year as the last prerequisite unit or in a later year.

Option 1: The follow-on alternative must be installed at least one year after the last prerequisite unit is installed.

The maximum lag year option specifies the maximum number of years between the installation of the last prerequisite unit and the installation of the follow-on alternative. The options are:

Option 0: The follow-on alternative must be installed in the same year as the last prerequisite unit.

Option 1: The follow-on alternative may be installed either in the same year as the last prerequisite unit or one year later.

Option -1: The follow-on alternative may be installed in the same year as the last prerequisite unit or anytime thereafter (no restriction).

During state evaluation the DP checks whether there are any feasible transitions from states in the previous year. If the previous-year state has more resources of any type than the current-year state, then the transition is not feasible. When an alternative has a prerequisite and the lag year option is specified, then the prerequisite resources must be installed in or before the previous year. Thus, the transition from a previous-year state is feasible only if the previous-year state has all the lag-year prerequisites for the current-year state.

To illustrate a minimum lag of one year, assume that the lag-year option was specified in the example given above; thus, two units of the first alternative must be installed at least one year before each unit of the second alternative type. Assume that the feasible states in the previous year are:

1	(0,0)	4	(3,0)	7	(3,1)
2	(1,0)	5	(4,0)	8	(4,1)
3	(2,0)	6	(2,1)	9	(4,2)

For current-year state (3,1), there would be feasible transitions from previous-year states 3, 4, 6, and 7. States 1 and 2 would not be allowed because they do not have the 2 units of the first alternative type required as prerequisites for the second alternative type unit. States 5, 8, and 9 would not be allowed because they have more units of the first alternative type (4) than are present in the current-year state (3). For current-year state (4, 2), there would be feasible

transitions from previous-year states 5, 8, and 9 – the only states with the 4 units of the first type that are required as prerequisites.

Another input parameter is the requirement option, which specifies whether a follow-on alternative must be installed if the prerequisite conditions are satisfied. When the requirement option is turned on, then the follow-on alternative must be installed by the last study year, except where the prerequisite conditions are first satisfied in the last study year and the minimum lag year option forces a one-year delay.

Still another input option is the retirement option, which specifies whether the prerequisite alternative should be retired when the follow-on alternative is installed. When the option is turned on, all prerequisite alternative units are retired when the follow-on unit is installed, or if one prerequisite unit enables the installation of multiple follow-on units, the prerequisite unit is retired when the first follow-on unit is installed. When the option is turned off, the prerequisite unit(s) continue to operate after the follow-on alternative is installed.

East planning alternative can reference up to four prerequisite alternatives. The lag year, requirement, and retirement options can be specified separately for each prerequisite. When there are multiple prerequisites, a follow-on alternative is installed only when all of the prerequisite conditions have been satisfied. If the prerequisite conditions are conflicting, then no units are installed.

With so many options, there is considerable flexibility in modeling different types of decisions. Examples include the phased installation of a new unit, simultaneous installation of multiple alternatives, sequencing the installation of multiple options, and timing of intervals between separate actions. Note that the lag year options allow for only a one-year interval. To establish a multi-year interval, you must use one or more dummy alternatives – alternatives that serve as follow-on units to the early stages and prerequisites to the later stages. In general, to establish a lag of N years requires N-1 dummy alternatives. The DP methodology precludes the specification of a multi-year lag without separate entities to represent each year of the lag. The DP also imposes restrictions on several specific combinations of lag year, requirement, and retirement options. Please refer to the information on planning alternatives in the EGEAS User's Guide for more specific information on these restrictions.

Following is an example that illustrates many of the features of dependency constraints:

Planning Alternative	Prerequisite Alternative	Retirement Option	Minimum Lag Option
CT	–	–	–
Dummy 1	CT (2)	No	1
Dummy 2	Dummy 1 (1)	No	1
CC	Dummy 2 (1)	No	1
	CT	Yes	0

This example involves the phased installation of a combined cycle unit. The first phase involves installation of two combustion turbines (alternative CT), either in the same year or in different years. The second phase is the conversion of the plant to combined cycle operation (alternative CC), at least 3 years after completion of the second CT. Because a lag of the 3-year interval, two dummy alternatives (Dummy 1 and Dummy 2) are needed. The specifics are as follows:

Alternative CT simply installs a combustion turbine, which will continue to operate until it reaches the end of its operating life or is forced into retirement by a later alternative.

Alternative Dummy 1 cannot be installed until at least one year after the second of two CT alternatives has been installed. This alternative has no purpose other than lengthening the interval between the phases; after the Dummy 1 alternative is installed, the CT alternatives continue to operate.

Alternative Dummy 2 cannot be installed until at least one year after Dummy 1, or more importantly, at least two years after the second CT alternative. As with Dummy 1, Dummy 2 has no effect on the operation of the CT alternatives.

Alternative CC represents conversion of the new unit to combined cycle operation. There are two prerequisite alternatives. The reference to the Dummy 2 alternative, with a minimum lag of one year, results in an overall interval of three years between the installation of the second CT and the conversion to combined cycle operation. The reference to the CT alternatives causes those original alternatives to be retired. The reason for retiring the CTs is that the CC alternative models the operation of the entire unit, including the combustion turbine portions.

Without the requirement option, the DP would be free to install the CT alternatives and never convert the unit to combined cycle operation. If you wanted to require that the plant be completed if it was started, you would turn on the requirement option associated with each of the prerequisite conditions. Note that forcing an alternative to retire does not affect its capital costs. Thus, the CT alternatives should include the capital costs associated with the combustion turbine portions of the plant and the CC alternative should include the capital costs associated with converting the plant to combined cycle operation. The operating characteristics of the CC alternative should reflect the operation of the entire unit because the CT alternatives are retired when it is installed.

Note that you could allow the CT alternatives to continue operating after the CC alternative was installed. In this case, the operating characteristics of the CC alternative would exclude the combustion turbine portions of the plant. However, in this case, the different portions of the plant would be treated as independent units, with possible conflicts in their operation (e.g., maintenance and forced outages). Thus, the approach illustrated with the example is recommended.

Note that two different planning alternatives can have the same prerequisites. In this case the number of prerequisite resources needed is the maximum of the number required by each of the selected alternative types (the number of prerequisites is not additive). During state generation, the DP performs prerequisite resource checking independently for each alternative type. During state evaluation, lag-year prerequisites for feasible transitions are also checked independently for each alternative type. When the retirement option is specified, prerequisites are retired by setting the retirement date to the end of the previous year (if the unit is not already retired). A resource may already be retired if:

the resource has already reached the end of its operating life, or

the resource was forced into retirement by a different planning alternative type.

Retirement of an Existing Resource

To model the retirement of an existing resource, you must specify a planning alternative that references the basic plant type of the resource to be retired. During state evaluation, the DP

calculates reserve margin, loss of load probability, and production cost. If the state contains an alternative that represents the retirement of an existing resource, then the DP sets the retirement date to the end of the previous year (if the resource is not already retired) so the resource will be excluded from the reliability and cost calculations. A resource may already be retired if:

the resource has already reached the end of its operating life,

the resource was forced into retirement by a different planning alternative type, or

the resource was forced into retirement by a previous resource of the same alternative type. (If there is more than one resource of an alternative that simply retires an existing resource, only the first resource has any effect).

Replacement of an Existing Unit

To model the replacement of an existing resource, you must specify a planning alternative that references both the basic plant type of an existing resource to be retired and the basic plant type of a generic unit to be installed in its place. The DP handles the existing resource retirement as described in the previous section of this document.

Mutually-Exclusive Alternatives

There are times when two or more planning options cannot be implemented together. For example, there might be several options for reducing the emissions by a particular plant, including: (1) switch to a low-sulfur coal, (2) add scrubbers, (3) convert the unit to burn gas, and (4) retire the unit. The DP optimization can be used to determine which of these strategies is the most economical, but you would want only one option selected in a particular expansion plan. Thus, EGEAS allows you to specify that a group of two or more planning alternatives are mutually-exclusive. In such a case, the DP can generate states that have any of the alternatives from the group, but it will never generate a state that has more than one alternative from the group. The same requirement applies to the optimal and suboptimal expansion plans: all of the alternatives may be included in the plans, but no plan will have more than one alternative from the group. You can specify several different groups of mutually-exclusive alternatives.

Single-Year Alternatives

A state is defined by the number of resources of each planning alternative. For standard alternatives, this number is the cumulative number of resources installed up through the current study year. For a single-year alternative, this number is the number of resources installed in the current year. Because of this distinction, the DP processes single-year alternatives differently in several ways. First the installation year for a single-year alternative is always the current year. Second, when the DP checks for feasible transitions from states in the previous year, it ignores the single-year alternatives. (The number of resources installed in the current year can be less than, greater than, or equal to the number installed in the previous year.) Third, the DP sets the retirement year of a single-year alternative equal to the installation year (units are installed on January 1 and retired on December 31), regardless of the operating life of the referenced basic plant type.

Multiple-Unit Capital Costs

During state evaluation, the DP calculates the committed cost and determines a backward pointer for each state. Included in the committed cost is the capital costs for all the installed planning alternatives. When you have supplied two different capital costs for a planning alternative, the

DP must decide which cost to use for each resource. In the first study year, the decision is straightforward: if the number of resources installed is NUNIT and the multiple-unit capital cost option is N, then the number of resources using the first (N1) and second (N2) capital costs are:

$$N1 = (NUNIT + N - 1) \quad / \quad N$$

$$N2 = NUNIT = N1$$

Note that integer division is used (e.g. 6/3 = 2, 7/3 = 2, 8/3 = 2).

In a later study year, the DP must account for resources installed in previous years. The DP determines the total number of resources using each cost that were installed in or before the current study year (as shown above), then subtracts the number of resources using each cost that were installed in or before the previous study year. If the cumulative number of resources for the current year is NUNITC, the cumulative number of resources for the previous year is NUNITX, and the multiple-unit capital cost option is N, then the number of resources installed in the current year that use the first (N1) and second (N2) capital costs are:

$$N1 = (NUNITC + N - 1)/N - (NUNITX + N - 1)/N$$

$$N2 = (NUNITC - NUNITX) - N1$$

Note again that integer division is used.

In the case of single-year alternatives, the DP considers only the number of resources installed in each year. Thus, the calculations of N1 and N2 are as shown above the first study year. In essence the DP starts a new group of resources each year rather than completing a partial group from the previous year.

Cost Postprocessor

The cost postprocessor is a tool for evaluating plans in terms of economic feasibility. When specified, the postprocessor evaluates every optimal and suboptimal expansion plan generated by the optimization. For this evaluation, the postprocessor calculates the interest coverage ratio, earning assets ratio, and percentage increase in system average rate for each year of the study period; it also calculates multi-year averages for these economic feasibility measures. Once calculated, the feasibility measures are compared with limiting values supplied. If any measure is outside the specified bound in any year or group of years, the plan is marked infeasible. EGEAS lets you decide whether infeasible plans should be simulated and reported or discarded.

Limitations

Because of limitations inherent in the methodology, there are several situations where the DP cannot guarantee an optimum solution. Several of these limitations involve annual changes in planning alternative characteristics. These limitations apply when an alternative can be installed in more than one year. If the alternative is limited to installation in only one year or if the alternative is a single-year type alternative, the limitations do not apply. Following is a list of cases where planning alternative characteristics may prevent the DP from determining an optimum solution:

A planning alternative has a short operating life, such that it could retire during the study period.

A planning alternative has characteristics that change from year to year, based on the age of the unit. Examples of such characteristics include rated, reserve, emergency, or operating capacity; forced outage rate; maintenance requirements; heat rate; hydro or storage energy limit. Note if a change in characteristics is tied to particular calendar years rather than the age of the unit, then the change is not a problem.

A planning alternative has a detailed cost modeled as variable O&M with an expenditure pattern.

A planning alternative has a detailed cost modeled as fixed O&M with an expenditure pattern and the possibility of being forced into retirement by another alternative. Note that optimality can be lost even if the alternative is not actually forced into retirement; just the potential for being retired is enough to affect optimality.

Another category of limitations applies only when the objective function is the levelized system average rate. The specific situations where optimality is not guaranteed are as follows:

The objective function is levelized system average rate and there is an existing or committed DSM program whose output can change from year to year.

The objective function is levelized system average rate and there is at least one DSM planning alternative.

The objective function is levelized system average rate and price elasticity is modeled (results in changes in system demand).

Practical Guidelines

For large optimization problems, the dynamic program can require a significant amount of computer resources – both time and storage space. Thus, you should formulate a problem carefully and select appropriate options. The basic objective is to obtain an optimal solution with the least amount of computer resources. Users modeling large utilities with many planning alternatives on small computers should be particularly careful. While methods exist for reducing the size of a problem, they can result in loss of the optimal solution. Thus, you need to strike a reasonable balance between accurate modeling and computer resource usage.

Following are some methods for reducing the computer resources for solving a DP optimization:

Reduce length of study period: You can achieve significant savings by shortening the study period. If the end of the study period is moved to an earlier year, some of the effect of those later years can still be captured, though less precisely, with the extension period. If no planning alternatives are needed in the early years, the study period can be changed to start when new resources are needed.

Reduce number of planning alternatives: Fewer alternatives produce fewer possible states; hence, state generation and evaluation take less time. Sometimes, you can eliminate some alternatives with the screening curve or some other screening tool. Another possibility is to divide the available alternatives into several smaller groups and optimize within each group, then optimize using the best members of each small group.

Tighten constraints: If not properly constrained, the DP may have to evaluate many states in each year that have no chance of being part of the optimal solution. Sometimes these uneconomic possibilities can be eliminated quickly with tighter reserve margin or tunnel constraints. Reducing the maximum number of superfluous units can significantly reduce the number of states.

Reduce production costing detail: A large part of the time needed for state evaluation is spent in production costing. Production costing can be speeded up by reducing the number of subperiods, by aggregating storage units into fewer storage projects, by eliminating the modeling of loading blocks, spinning reserve, or economy interchange, and by combining small units.

The danger in reducing the size of a problem is that you may eliminate the optimum solution. Following is a list of possible causes for loss of optimality:

Too many feasible states generated: If the DP reaches the upper limit on the number of feasible states generated, then it must stop generating states. The optimal expansion plan may include a state that could not be generated. The DP prints a warning message to you and proceeds to the next year.

Too many feasible states to be retained: If the DP generates more states in a year than can be saved for the optimization, then the DP must choose which states to retain and which to discard. You can specify several different methods for selecting the plans to retain, but none of the methods guarantees that the optimal expansion plan is preserved. The DP lists which states are retained and which are discarded.

Overly-restrictive constraints: If you tighten constraints for the sole purpose of reducing the computational requirements, then the optimal solution to the original problem may be excluded. When the optimal solution determined by the DP is at one of the constraint boundaries, you should investigate whether a better solution could be obtained with a looser constraint.

Superfluous units: In some cases, the optimal solution to a problem may have extra resources in some years. It may be cheaper to install a resource a year earlier than needed, or the operational savings might outweigh the capital costs for extra resources, or restriction on installation years might require installation before a resource's generation is needed.

7
FINANCIAL ANALYSIS

Overview

EGEAS was designed to provide basic financial information and also provide detailed cost information if you want to complete a more detailed financial analysis in a financial planning tool such as FINANCE+. When providing input to EGEAS, you must choose between three cost analysis formats. These formats determine how EGEAS calculates capital costs, what input parameters are required, and what modeling capabilities are available. The three formats are:

Format 1 – Levelized fixed charges: Capital fixed charges are calculated using an installation cost and a levelized carrying charge. The same levelized value is then applied in each year of a resource's book life. Construction costs are not modeled and economic feasibility measures are not calculated.

Format 2 – Annual fixed charges: Annual (non-levelized) fixed charges are calculated for each year of a resource's book life using an installation cost, book and tax depreciation, and system rates of return on rate base. Construction costs are not modeled and economic feasibility measures are not calculated.

Format 3 – Construction costs: Construction costs are modeled using direct annual construction expenditures and AFUDC rates. Annual (non-levelized) fixed charges are calculated for each year of the resource's book life using the resource's gross book value, book and tax depreciation, and system rates of return on rate base. If any portion of CWIP is allowed in the rate base, fixed charges are calculated for the construction period as well. Economic feasibility indicators can be calculated.

Format 1 – Levelized Fixed Charges

Cost analysis format 1 is the simplest of the three formats, in terms of calculations and input requirements. With format 1, you must supply a capital installation cost (in base year dollars per kilowatt of capacity), a levelized carrying charge (in percent), and a book life (in years) for each new resource. You also specify a set of escalation rates for projecting the cost to the installation year. The installation cost, or gross book value, includes direct construction expenditures and interest during construction, also known as allowance for funds used during construction (AFUDC). The levelized carrying charge incorporates depreciation, interest payments, taxes, and return on equity, levelized over the book life of the resource. This carrying charge is applied to the installation cost to yield the uniform annual cost of owning a facility.

For a generic resource, you can specify two different capital installation costs and a multi-unit capital cost option. If the option is set to N, then resources are assumed to be installed in groups of N units. The first unit in each group uses the first capital cost and the remaining units use the second cost. As an example, assume that the multi-unit capital cost option is 4. Then the first capital cost would be used for units 1, 5, 9, 13, etc. and the second cost would be used for the remaining units (2, 3, 4, 6, 7, 8, 10, 11, 12, 14, etc.). Thus, you can assign a higher cost to the first unit at a site to represent the costs of site preparation or shared facilities (e.g. transmission lines or fuel handling equipment). The Benders' decomposition and screening curve analysis

methodologies cannot distinguish multiple resources (the number of resources may be fractional), so the weighted average of the two costs within a complete group is used for every resource (or fraction of a resource).

With format 1, most of the effort involved in calculating capital costs is concentrated outside of EGEAS. Items such as AFUDC, depreciation, taxes, interest payments, and annual fixed charges must be calculated and then used to determine the installation cost and levelized carrying charge required by EGEAS. The advantage of this method is that you can tailor the calculations as needed for a specific utility. The disadvantage is that EGEAS can neither report annual regulatory carrying charges nor calculate economic feasibility measures.

Calculations

EGEAS calculates a levelized fixed charge value for each new resource. This value is then applied in every year of the resource's book life that falls within the study or extension period. Note that although the fixed charges remain constant in current year dollars, the present worth of the fixed charges declines as the resource ages because the charges are discounted back a greater number of years. The levelized carrying charge must be input.

Define:

IDBSYR	=	base year of data base
IY	=	installation year
BPRCAP	=	rated capacity (MW)
BPCPZ1	=	capital installation cost ($/KW) Note: BPCPZ1 represents the cost of the first unit within a group. For the costs of the remaining units in the group, simply replace BPCPZ1 by BPCPZ2.
BPLEVC	=	levelized carrying charge (percent)
UKZFIX	=	capital fixed charges ($)

Determine gross book value (in installation year $):

CZ	=	BPCPZ1, escalated from base year IDBSYR to installation year IY
ZINSTL	=	CZ * (BPRCAP * 1000)

Calculate levelized fixed charges (in current year $):

UKZFIX = ZINSTL * (BPLEVC/100)

Format 2 – Annual Fixed Charges

Cost analysis format 2 is similar to format 1 except that EGEAS calculates annual rather than levelized fixed charges. As with format 1, you must supply a capital installation cost (in base year dollars per kilowatt of capacity), a set of escalation rates for projecting the cost to the installation year, and a book life (in years) for each new resource. However, format 2 does not use a levelized carrying charge. Instead, you must specify the AFUDC amount as a fraction of the installation cost, and the fraction of that AFUDC financed by debt (as opposed to equity). Tax depreciation schedules are needed for each resource or class of resources. You also must specify the capital structure (percent debt, common stock, preferred stock), debt interest rate,

allowed rate of return on equity, cost of preferred stock, income tax rate, and property tax rate for the system as a whole.

EGEAS calculates the capital fixed charges for each year of a resource's book life as the sum of book depreciation, AFUDC adjustment, return on rate base, and property taxes. You can choose between two variations of capital cost calculations. With the standard method, the return after the installation is based on the mid-year value of rate base. With the alternate method, the return is based on the beginning of year value of rate base.

Pre-tax return on investment: EGEAS calculates a pre-tax rate of return on rate base (RORB) using the capital structure, the debt interest rate, the rate of return allowed on equity, the cost of preferred stock, and the tax rate. For each year of a resource's book life, EGEAS applies the RORB to the mid-year value of rate base (standard method) or the beginning of year value of rate base (alternate method) to obtain the pre-tax return on investment.

Book depreciation: EGEAS calculates the straight-line book depreciation by dividing the installation cost by the book life of the resource.

AFUDC adjustment: For this term, EGEAS computes the straight-line depreciation of the equity portion of the AFUDC. This value is then multiplied by the tax rate divided by one minus the tax rate.

Property taxes: For each year of the resource's book life, EGEAS calculates property taxes by applying the user-supplied property tax rate to the mid-year rate base value (standard method) or to the beginning of year rate base value (alternate method).

With format 2, calculations involving annual construction expenditures, AFUDC accrual, and construction work in progress (CWIP) balances must be performed outside of EGEAS. You supply the results of these external calculations and EGEAS then determines the investment return history with annual fixed charges. The advantages of this method are that you can tailor the AFUDC calculations as needed for a specific utility and that EGEAS can report annual regulatory carrying charges. The disadvantage is that EGEAS cannot calculate economic feasibility measures.

Capital Cost Structure by Resource

You can specify multiple sets of return on rate base parameters for use with different resources. Each set includes the capital structure (percent common stock, preferred stock, debt), rates of return (return on equity, cost of preferred stock, debt interest rate), and tax rates (income and property). For each resource, you can specify which capital structure data set to use for capital cost calculations. This facilitates the evaluation of independent power producer resources or DSM resources that would have a different capital cost structure than the utility.

Rate of Return on Rate Base

The calculation of annual fixed charges uses the rate of return on rate base. For each year y, EGEAS calculates a rate of return based on the capital structure, allowed rate of return on equity, debt interest rate, and tax rate supplied.

Define:

ZRCOMM(y)	=	percent common stock in capital structure
ZRPREF(y)	=	percent preferred stock in capital structure
ZRDEBT(y)	=	percent debt in capital structure
ZREINT(y)	=	rate of return allowed on equity (percent)
ZRPINT(y)	=	cost of preferred stock (percent)
ZRDINT(y)	=	debt interest rate (percent)
ZRTAXR(y)	=	annual effective state and federal tax rate (percent)
ZRETRB(y)	=	system rate of return on rate base (percent)

Calculate rate of return on rate base (percent):

$$ZRETRB = \frac{ZRCOMM * ZREINT + ZRPREF * ZRPINT}{100 - ZRTAXR} + \frac{ZRDEBT * ZRDINT}{100}$$

Note that the index y is omitted for simplicity in this formula.

Fixed Charges Calculations

EGEAS calculates annual fixed charges values for each new resource. Starting with the installation cost of the resource, the program calculates a declining series of fixed charge values which depend on the cost of capital, book and tax depreciation, and tax rates. In addition to supplying the installation cost, you must also specify equity and debt AFUDC as percentages of this cost.

Define:

IDBSYR	=	base year of data base
IY	=	installation year
BPRCAP	=	rated capacity (MW)
IBKLBP	=	book life (years)
BPCPZ1	=	capital installation cost ($/KW) Note: BPCPZ1 represents the cost of the first unit within a group. For the costs of the remaining units in the group, simply replace BPCPZ1 by BPCPZ2.
BPAFCP	=	default AFUDC portion of gross book value (percent)
BPAFCD	=	default debt portion of AFUDC (percent)
ZTFRAC(n)	=	tax depreciation for nth year of resource's tax life (percent) Note: If no tax depreciation schedule is supplied for a resource, then straight-line book depreciation is used and income tax is not deferred.
ZRTAXR(y)	=	income tax rate for year y (percent)
ZRPROP(y)	=	property tax rate for year y (percent)
ZRETRB(y)	=	system rate of return on rate base for year y (percent)

UKDEPB(y) = book depreciation for year y ($)

UKDEPA(y) = AFUDC depreciation for year y ($)

UKDEPT(y) = tax depreciation for year y ($)

UKDITX(y) = deferred income tax for year y ($)

UKADIT(y) = accumulated deferred income tax for year y ($)

UKZNBV(y) = net book value for end of year y ($)

UKZRBV(y) = rate base value for end of year y ($)

UKZFIX(y) = capital fixed charges for year y ($)

Determine gross book value (in installation year $):

CZ = BPCPZ1, escalated from base year IDBSYR to installation year IY

ZINSTL = CZ * (BPRCAP * 1000)

Break down gross book value ($) into debt AFUDC (UKCWAD), equity AFUDC (UKCWAE), and direct cost (UKCWDN) portions:

CWAT = ZINSTL * (BPAFCP / 100)

UKCWAD = CWAT * (BPAFCD / 100)

UKCWAE = CWAT - UKCWAD

UKCWDN = ZINSTL - CWAT

Initialize balances at end of year immediately preceding the installation year (y_0 = IY – 1):

Net book value and rate base value:

UKZNBV(y_0) = ZINSTL

UKZRBV(y_0) = ZINSTL

Accumulated deferred income tax:

UKADIT(y_0) = 0

The following three quantities are calculated once. They have the same value in every year y of the resource's book life:

Book depreciation ($)

UKDEPB = ZINSTL/IBLKBP

AFUCD depreciation ($)

UKDEPA = UKCWAE/IBKLBP

Book depreciation excluding AFUDC ($)

DEPD = UKDEPB - UKDEPA

The remaining calculations are performed for every year y of the resource's book life to obtain the annual fixed charges.

Tax depreciation, using input tax depreciation schedule ($):

During the tax life of the resource:

$$UKDEPT(y) = (ZINSTL - UKCWAE) * \frac{ZTFRAC(y - IY + 1)}{100}$$

After the tax life of the resource:

UKDEPT(y) = 0

No tax depreciation schedule supplied:

UKDEPT(y) = DEPD

Deferred income tax ($):

$$UKDITX(y) = (UKDEPT(y) - DEPD) * \frac{ZRTAXR(y)}{100}$$

Accumulated deferred income tax ($):

UKADIT(y) = UKADIT(y-1) + UKDITX(y)

Net book value ($):

UKZNBV(y) = UKZNBV(y-1) - UKDEPB(y)

Rate base value ($):

UKZRBV(y) = UKZNBV(y) - UKADIT(y)

Mid-year rate base value ($):

$$RBASM = \frac{UKZRBV(y - 1) + UKZRBV(y)}{2}$$

Annual fixed charges ($):

Standard method (based on the mid-year value of rate base):

$$UKZFIX(y) = RBASM * \frac{ZRETRB(y) + ZRPROP(y)}{100} + UKDEPB + UKDEPA * \frac{ZRTAXR(y)}{100 - ZRTAXR(y)}$$

Alternate method (based on the beginning of year value of rate base):

$$UKZFIX(y) = UKZRBV(y - 1) * \frac{ZRETRB(y) + ZRPROP(y)}{100} + UKDEPB + UKDEPA * \frac{ZRTAXR(y)}{100 - ZRTAXR(y)}$$

Format 3 – Construction Costs

Cost analysis format 3 is the most detailed of the three methods and requires the most input data. There are two stages of calculations: construction cost modeling and annual fixed charges calculation. For the first stage, you must supply direct construction expenditures for each year of a resource's construction period. EGEAS computes the amount of AFUDC accrued each year and maintains construction work in progress (CWIP) balances, broken down into AFUDC and direct cost portions. You must supply a system-wide AFUDC rate and must specify whether AFUDC interest is applied to the total CWIP balance (compound interest option) or just to the direct costs (simple interest option). You may also specify the percentage of the direct

construction expenditures allowed in rate base each year. When CWIP is allowed in rate base, EGEAS calculates the fixed charges associated with a return on rate base for those assets; however, the assets are not depreciated until the resource goes on-line. The CWIP balance at the end of the last construction period year (including any costs allowed in rate base) is the gross book value of the resource. This gross book value corresponds to the installation cost input for cost analysis formats one and two.

The second stage of calculations, to determine the annual fixed charges, is basically the same as for cost analysis format 2. Construction of a new resource starts at least one year before the resources installation year. Thus, the construction period for a resource installed in an early year of the study period often starts before the beginning of the study period. On the other hand, the book life of a new resource often extends beyond the end of the study period. EGEAS performs calculations over the entire life cycle of the resource, but only the fixed charges that fall within the study and extension periods are included in the cost of an expansion plan.

Because cost parameters (e.g., AFUDC rates and capital structure) are specified only for years in or after the base year, EGEAS cannot perform complete cost calculations before the base year. However, construction of a committed resource may have started before that date. Thus, for a committed resource, you may supply starting values for AFUDC and CWIP balances as of the start of the base year. If you do choose to specify a construction period that starts before the base year, EGEAS performs the calculations with several assumptions. First, the AFUDC rate and capital structure are set to the base year values. Second, there is no cost escalation or discounting.

With cost analysis format 3, more of the capital cost calculations are performed by EGEAS than is the case with formats 1 and 2. Correspondingly, the input requirements are more extensive as well. The advantages of this method are that all the capital cost calculations are based on consistent economic assumptions, that annual regulatory carrying charges are reported, and that economic feasibility measures can be computed. The disadvantages are that the data requirements are greater and you have less ability to tailor the calculations to the needs of a specific utility.

You can choose between two variations in the calculations of capital costs when you have an EGEAS data base with the construction cost format. The differences between the two methods are as follows:

With the standard method, direct construction costs that are added to the rate base do not accumulate AFUDC. With the alternate method, AFUDC is accumulated on these costs. The equations used for both methods to compute the amount of AFUDC accrued each year are documented later in this section.

With the standard method, the fixed charges during the construction period include a return on only the direct construction cost that is allowed in rate base. With the alternate method, the fixed charges include a return on the AFUDC as well as the direct cost in rate base, but the fixed charges are reduced by the annual AFUDC accrual. The equations used for both methods are given later in this section.

With the standard method, the return after installation is based on the mid-year value of rate base. With the alternate method, the return is based on the beginning of year value of rate base. For the equations used for both methods in calculating the annual fixed charges during the resource's book life, see the **Format 2 – Annual Fixed Charges** section.

Separating Equity and Debt Portions of AFUDC

The calculation of annual fixed charges requires that the equity and debt portions of AFUDC be kept separately. The calculations, based on the capital structure and rates of return, determine how the AFUDC accrual each year (y) should be split into equity and debt portions.

Define:

y	=	year of study period
ZRCOMM(y)	=	percent common stock in capital structure
ZRPREF(y)	=	percent preferred stock in capital structure
ZRDEBT(y)	=	percent debt in capital structure
ZREINT(y)	=	rate of return allowed on equity (percent)
ZRPINT(y)	=	cost of preferred stock (percent)
ZRDINT(y)	=	debt interest rate (percent)
DFRAC(y)	=	debt fraction of AFUDC accrual

Debt fraction of AFUDC accrual:

$$DFRAC = \frac{ZRDEBT * ZRDINT}{ZRCOMM * ZREINT + ZRPREF * ZRPINT + ZRDEBT * ZRDINT}$$

Note that the index y is omitted for simplicity.

Construction Cost Calculations

When the cost analysis format specifies construction cost modeling, EGEAS starts with direct construction cost expenditures and calculates the allowance for funds used during construction (AFUDC) accrual and the construction work in progress (CWIP) balance for each year of the construction period. When a portion of CWIP is allowed in the rate base, then EGEAS calculates the fixed charges during the construction period as well. For committed resources whose construction started before the end of the base year, EGEAS accepts starting values for AFUDC and CWIP balances. The CWIP balances at the end of the construction period are saved for calculation of fixed charges during the book life of the resource. These book life calculations are documented in the **Format 2 – Annual Fixed Charges** section. The equations for construction cost calculations are included here.

Define:

IDBSYR	=	base year of data base
IY	=	installation year
BPRCAP	=	rated capacity (MW)
IBKLBP	=	book life (years)
BPDCZ1	=	direct (overnight) construction cost (to-go) ($/KW) Note: BPDCZ1 represents the cost of the first unit within a group. For the costs of the remaining units in the group, simply replace BPDCZ1 by BPDCZ2.

NYRSZC	=	number of years in construction period Note: If no construction schedule is specified for a resource, the resource is assumed to be constructed in less than a year with no AFUDC. The gross book value is the same as the direct construction cost
ZCFRAC(j)	=	percentage of construction cost expended in the year before installation
BPCWPR	=	percent of CWIP allowed in rate base
BPCWP1	=	CWIP balance (including AFUDC) at the start of base year ($/KW) (for committed resources whose construction began before the base year)
BPAFE1	=	equity AFUDC balance at start of base year ($/kW) (for committed resources whose construction began before the base year)
BPAFD1	=	debt AFUDC balance at start of base year ($/kW) (for committed resources whose construction began before the base year)
IAFCZA	=	AFUDC compounding option (1=compound, 2=simple)
y	=	year of study period
ZAFCRT(y)	=	AFUDC rate (percent)
ZRETRB(y)	=	system rate of return on rate base (percent)
UKDIRR(y)	=	current year direct cost added to rate base ($)
UKDIRN(y)	=	current year direct cost not added to rate base ($)
UKAFCE(y)	=	equity AFUDC accrual ($)
UKAFCD(y)	=	debt AFUDC accrual ($)
UKCWDR(y)	=	year-end CWIP balance – direct cost in rate base ($)
UKCWDN(y)	=	year-end CWIP balance – direct cost not in rate base ($)
UKCWAE(y)	=	year-end CWIP balance – equity AFUDC portion ($)
UKCWAD(y)	=	year-end CWIP balance – debt AFUDC portion ($)
UKDEPB(y)	=	book depreciation ($)
UKDEPA(y)	=	AFUDC depreciation ($)
UKDEPT(y)	=	tax depreciation ($)
UKDITX(y)	=	deferred income tax ($)
UKADIT(y)	=	accumulated deferred income tax ($)
UKZNBV(y)	=	net book value for end of year y ($)
UKZRBV(y)	=	rate base value for end of year y ($)
UKZFIX(y)	=	capital fixed charges for year y ($)

Calculate direct construction cost to-go (base year $):

ZCSTRC = BPDCZ1 * (BPRCAP * 1,000)

Determine construction period:

First year = y_1 = IY - NYRSZC

Last year = y_n = IY - 1

Initialize CWIP balances at end of year, for $y_0 = y_1 - 1$:

For resources whose construction period starts in or after the base year, initialize CWIP balances to zero:

UKCWDR(y0) = 0

UKCWDN(y0) = 0

UKCWAD(y0) = 0

UKCWAE(y0) = 0

For committed or existing resources, costs incurred before the base year are represented by the initial values of the CWIP (BPCPW1) and the debt and equity AFUDC portions of CWIP (BPAFD1, BPAFE1). These initial values are balances as of the start of the base year:

Resource direct cost - total (CWIPD)

- portion in rate base (UKCWDR)
- portion not in rate base (UKCWDN)

CWIPD = (BPCWP1 – BPAFD1 – BPAFE1) * (BPRCAP * 1000)

$$\text{UKCWDR(yo)} = \text{CWIPD} * \frac{\text{BPCWPR}}{100}$$

$$\text{UKCWDN(yo)} = \text{CWIPD} * \left(1 - \frac{\text{BPCWPR}}{100}\right)$$

Resource AFUDC - equity portion (UKCWAE)

- debt portion (UKCWAD)

UKCWAE(y_0) = BPAFE1 * (BPRCAP * 1,000)

UKCWAD(y_0) = BPAFD1 * (BPRCAP * 1,000)

The remaining calculations are performed for every year y of the resource's construction period (y_1 to y_n):

Current year direct cost - total (DIRT)

- portion added to rate base (UKDIRR)
- portion not added to rate base (UKDIRN)

DIRT = ZCSTRC * ZCFRAC(IY - y), escalated to year y

$$UKDIRR(y) = DIRT * \frac{BPCWPR}{100}$$

$$UKDIRN(y) = DIRT - UKDIRR(y)$$

Current year AFUDC accrual - total

Standard Method:

with compounding (IAFCZA = 1)

$$AFCT = ZAFCRT(y) * \left(\frac{UKDIRN(y)}{2} + UKCWDN(y-1) + UKCWAE(y-1) + UKCWAD(y-1) \right)$$

without compounding (IAFCZA = 2)

$$AFCT = ZAFCRT(y) * \left(\frac{UKDIRN(y)}{2} + UKCWDN(y-1) \right)$$

Alternate Method:

with compounding (IAFCZA = 1)

$$AFCT = ZAFCRT(y) * \left(\frac{UKDIRN(y) + UKDIRR(y)}{2} + UKCWDN(y-1) + UKCWDR(y-1) + UKCWAE(y-1) + UKCWAD(y-1) \right)$$

without compounding (IAFCZA = 2)

$$AFCT = ZAFCRT(y) * \left(\frac{UKDIRN(y) + UKDIRR(y)}{2} + UKCWDN(y-1) + UKCWDR(y-1) \right)$$

Current year AFUDC accrual - debt (UKAFCD)

- equity (UKAFCE)

$UKAFCD(y) = DFRAC * AFCT$

$UKAFCE(y) = AFCT - UKAFCD(y)$

Year-end CWIP balances - direct cost in rate base (UKCWDR)

- direct cost not in rate base (UKCWDN)

$UKCWDR(y) = UKCWDR(y-1) + UKDIRR(y)$

$UKCWDN(y) = UKCWDN(y-1) + UKDIRN(y)$

Year-end CWIP balances - equity AFUDC (UKCWAE)

- debt AFUDC (UKCWAD)

$$UKCWAE(y) = UKCWAE(y-1) + UKAFCE(y)$$

$$UKCWAD(y) = UKCWAD(y-1) + UKAFCD(y)$$

No depreciation during construction period:

$$UKDEPB(y) = 0$$

$$UKDEPA(y) = 0$$

$$UKDEPT(y) = 0$$

$$UKDITX(y) = 0$$

$$UKADIT(y) = 0$$

Year-end net book value in rate base ($):

before last year y_n

$$UKZNBV(y) = UKCWDR(y)$$

last year of construction period

$$UKZNBV(y) = UKCWDR(y) + UKCWDN(y) + UKCWAE(y) + UKCWAD(y)$$

Note: ZINSTL is the gross book value for the resource. It is used in calculating the annual fixed charges during the resource's book life. (See the **Format 2 - Annual Fixed Charges** section.)

Year-end rate base value (same as net book value ($):

$$UKZRBV(y) = UKZNBV(y)$$

Fixed Charges ($):

Standard Method:

first construction year (y_1)

$$UKZFIX(y) = \frac{ZRETRB(y)}{100} * \frac{UKCWDR(y)}{2}$$

not first construction year

$$UKZFIX(y) = \frac{ZRETRB(y)}{100} * \frac{UKCWDR(y-1) + UKCWDR(y)}{2}$$

Alternate method:

first construction year (y_1)

$$UKZFIX(y) = \frac{ZRETRB(y)}{100} * \frac{UKCWDR(y)}{2}$$

$$- \frac{UKAFCE(y) * 100}{100 - ZRTAXR(y)} - UKAFCD(y)$$

not first construction year

$$UKZFIX(y) = \frac{ZRETRB(y)}{100} * \left(\frac{UKDIRR(y)}{2} + UKCWDR(y-1) + UKCWAD(y-1) + UKCWAE(y-1) \right) - \frac{UKAFCE(y) * 100}{100 - ZRTAXR(y)} - UKAFCD(y)$$

Capital Expenditures Extended Beyond Installation Year

For existing resources as well as committed and generic resources, you can specify capital expenditures for capital improvements during a resource's operation. These capital expenditures extend from the resource's installation year forward to a maximum of 30 years.

You have the option of specifying capital expenditures in units of base-year dollars per kilowatt of rated capacity or as percentages of the direct construction cost. Expenditures can be either positive or negative. Also, you can specify a separate trajectory for escalating the expenditures from year to year. If no trajectory is supplied, EGEAS uses the trajectory specified for construction costs. There is no construction period and no AFUDC associated with these expenditures. The book life for the expenditures in a particular year is the number of years remaining until the end of the resource's basic book life. Straight-line depreciation is used for both book and tax purposes. If expenses are incurred after the resource's book life is over, the book life is set to one year and expenses are depreciated fully in the year they are incurred.

Sample Calculation

EGEAS Data Base – Input Parameters

General data:

Base year	1987	
Discount rate	8.14 %	
Capital structure	Common stock	38 %
	Preferred stock	12 %
	Debt	50 %
Rates of return	Common stock	10 %
	Preferred stock	7%
	Debt	7%
Income tax rate	34 %	
Property tax rate	2%	
AFUDC rate	7% with compounding	

Unit Data:

Capacity	250 MW	
Installation year	1994	
Operating life	5 years (1994 to 1998)	
Book life	5 years (1994 to 1998)	
Construction cost (1987$)	1400 $/kW	
Construction period	3 years (1991 to 1993)	
Construction expenditure pattern	1 year before installation (1993)	40 %
	2 years before installation (1992)	50 %
	3 years before installation (1991)	10%
	CWIP in rate base	0 %
Tax depreciation	year 1 (1994)	50 %
	year 2 (1995)	30%
	year 3 (1996)	10%
	year 4 (1997)	10%
Income tax rate	34 %	
Property tax rate	2%	
AFUDC rate	7% with compounding	

Preliminary Calculations

Rate of return on rate base:

$$\frac{(0.38)(0.10)+(0.12)(0.07)}{1.0-0.34}+(0.50)(0.07)=0.105303$$

Debt AFUDC as a fraction of total AFUDC accrual:

$$\frac{(0.50)(0.07)}{(0.38)(0.10)+(0.12)(0.07)+(0.50)(0.07)}=0.429975$$

Direct Construction Costs, in thousands of dollars (K$)

Total direct cost (1987$):

$$\frac{(250\,\text{MW})(1{,}000\,\text{kW/MW})(1{,}400\,\$/\text{kW})}{1{,}000\,\$/\text{K\$}}=350{,}000\,\text{K\$}$$

1991: (0.10) (350,000 K$) $(1.04)^4$ = 40,945.05 K$

1992: (0.50) (350,000 K$) $(1.04)^5$ = 212,914.26 K$

1993: (0.40) (350,000 K$) $(1.04)^6$ = <u>177,144.66 K$</u>

431,003.97 K$

AFUDC Calculations

1991: total (40,945.05)(.07/2) = 1,433.08

	debt	(1,433.08)(.429975)	=	616.19
	equity	1,433.08 - 616.19	=	816.89
	end-year CWIP = 40,945.05 + 1,433.08		=	42,378.13
1992:	total	(42,378.13)(.07) + (212,914.26)(.07/2)	=	10,418.47
	debt	(10,418.47)(.429975)	=	4,479.68
	equity	10,418.47 - 4,479.68	=	5,938.79
	end-year CWIP = 42,378.13 + 212,914.26 +10,418.47		=	265,710.86
1993:	total	(265,710.86)(.07) + (177,144.66)(.07/2)	=	24,799.98
	debt	(24,799.98)(.429975)	=	10,663.30
	equity	24,799.98 - 10,663.30	=	14,136.68
	end-year CWIP = 265,710.86 + 177,144.66 + 24,799.98		=	467,655.50

Gross Book Value

	Direct Costs	Equity AFUDC	Debt AFUDC	Total
1991	40,945.05	816.89	616.19	42,378.13
1992	212,914.26	5,938.79	4,479.68	223,332.73
1993	177,144.66	14,136.68	10,663.30	201,944.64
	431,003.97	**20,892.36**	**15,759.17**	**467,655.50**

Depreciation

Book depreciation:	(467,655.50)/5	=	93,531.10
AFUDC depreciation:	(20,892.36)/5	=	4,178.47
Book depreciation, excluding AFUDC		=	89,352.63

Fixed Charges - 1994

Tax depreciation:	(467,655.50 - 20,892.36)(.50)	=	223,381.57
Deferred income tax:	(223,381.57 - 89,352.63)(.34)	=	45,569.84
Accumulated deferred income tax:	=		45,569.84
Net book value:	467,655.50 - 93,531.10	=	374,124.40
Rate base value (end-year)	374,124.40 - 45,569.85	=	328.554.56
Rate base value (mid-year):	$\frac{467,655.50 + 328,554.56}{2}$	=	398,105.03

Fixed charges:	$(398{,}105.03)(0.105303 + 0.02) + 93{,}531.10 + \frac{(4{,}178.47)(0.34)}{(1.0 - 0.34)}$	=	145,567.40

Fixed Charges - 1995

Tax depreciation:	(467,655.50 - 20,892.36)(.30)	=	134,028.95
Deferred income tax:	(134,028.95 - 89,352.63)(.34)	=	15,189.95
Accum. deferred income tax:	45,569.84 + 15,189.95	=	60,759.79
Net book value:	374,124.40 - 93,531.10	=	280,593.30
Rate base value (end-year)	280,593.30 - 60,759.79	=	219,833.51
Rate base value (mid-year):	$\frac{328{,}554.56 + 219{,}833.51}{2}$	=	274,194.04
Fixed charges:	$(274{,}194.04)(0.105303 + 0.02) + 93{,}531.10 + \frac{(4{,}178.47)(0.34)}{(1.0 - 0.34)}$	=	130,040.98

Fixed Charges - 1996

Tax depreciation:	(467,655.50 - 20,892.36)(.10)	=	44,676.32
Deferred income tax:	(44,676.32 - 89,352.63)(.34)	=	-15,189.95
Accum. deferred income tax:	60,759.79 - 15,189.95	=	45,569.84
Net book value:	280,593.30 - 93,531.10	=	187,062.20
Rate base value (end-year)	187,062.23 - 45,569.84	=	141,492.39
Rate base value (mid-year):	$\frac{219{,}833.51 + 141{,}492.39}{2}$	=	180,662.95
Fixed charges:	$(180{,}662.95)(0.105303 + 0.02) + 93{,}531.10 + \frac{(4{,}178.47)(0.34)}{(1.0 - 0.34)}$	=	118,321.22

Fixed Charges - 1997

Tax depreciation:	(467,655.50 - 20,892.36)(.10)	=	44,676.32
Deferred income tax:	(44,676.32 - 89,352.63)(.34)	=	-15,189.95
Accum. deferred income tax:	45,569.84 - 15,189.95	=	30,379.89
Net book value:	187,062.20 - 93,531.10	=	93,531.10
Rate base value (end-year)	93,531.10 - 30,379.89	=	63,151.21
Rate base value (mid-year):	$\frac{141{,}492.39 + 63{,}151.21}{2}$	=	102,321.80

Fixed charges:	$(102,321.80)(0.105303 + 0.02) + 93,531.10 + \frac{(4,178.47)(0.34)}{(1.0-0.34)}$	=	108,504.87

Fixed Charges - 1998

Tax depreciation:	(467,655.50 - 20,892.36)(.00)	=	0.00
Deferred income tax:	(0.00 - 89,352.63)(0.34)	=	-30,379.89
Accum. deferred income tax:	30,379.89 - 30,379.89	=	0.00
Net book value:	93,531.10 - 93,531.10	=	0.00
Rate base value (end-year)	0.00 - 0.00	=	0.00
Rate base value (mid-year):	$\frac{63,151.21+0.00}{2}$	=	31,575.60
Fixed charges:	$(31,575.60)(0.105303 + 0.02) + 93,531.10 + \frac{(4,178.47)(0.34)}{(1.0-0.34)}$	=	99,640.16

Cumulative Present Worth of Fixed Charges

1994:	$145,567.40/(1.0814)^{7}$	=	84,170.43
1995:	$130,040.98/(1.0814)^{8}$	=	69,532.73
1996:	$118,321.22/(1.0814)^{9}$	=	58,503.97
1997:	$108,504.87/(1.0814)^{10}$	=	49,611.87
1998:	$99,640.16/(1.0814)^{11}$	=	42,129.32
			303,948.32

Replacement unit installed in 1999

Escalate present worth cost of original unit from 1994 to 1999, and discount back the additional 5 years, to keep it in base year (1987) dollars:

$$(303,948.32)\frac{1.04**(1999-1994)}{(1.0814)**(1999-1994)} = 250,054.47$$

For replacement units, levelized fixed charges are used. If A1 is the annual cost in current year dollars and C1 is the cumulative present worth cost, then:

$$C_1 = \frac{A1}{(1.0814)**(1999-1987)} + \frac{A1}{(1.0814)**(2000-1987)} + \frac{A1}{(1.0814)**(2001-1987)} + \frac{A1}{(1.0814)**(2002-1987)} + \frac{A1}{(1.0814)**(2003-1987)}$$

$$= \frac{A1}{(1.0814)**(1999-1987)} * \frac{1.0-(1.0/1.0814)**5}{1.0-(1.0/1.0814)}$$

Setting C_1 = 250,054.47 and solving for A_1 yields: A_1 = 148,668.03

In present worth dollars, the fixed charges for each year are:

1999:	$148,668.03/(1.0814)^{12}$	=	58,127.44
2000:	$148,668.03/(1.0814)^{13}$	=	53,752.03
2001:	$148,668.03/(1.0814)^{14}$	=	49,705.96
2002:	$148,668.03/(1.0814)^{15}$	=	45,964.45
2003:	$148,668.03/(1.0814)^{16}$	=	42,504.58

Replacement unit installed in 2004

Levelized annual cost:	$A_2 = A_1$ (1.04) ** (2004 - 1999)	=	180,877.39
Present worth fixed charges:			
2004:	$180,877.39/(1.0814)^{17}$	=	47,820.72
2005:	$180,877.39/(1.0814)^{18}$	=	44,221.12
2006:	$180,877.39/(1.0814)^{19}$	=	40,892.47
2007:	$180,877.39/(1.0814)^{20}$	=	37,814.38
2008:	$180,877.39/(1.0814)^{21}$	=	34,967.99

Extension Period Capital Costs (some examples)

Extension period - 1997 to 1999:	
1997 (original unit):	49,611.87
1998 (original unit):	42,129.32
1999 (replacement unit):	58,127.44
	149,868.63 K$
Extension period - 1999 to 2003:	
1999 (replacement unit):	58,127.44
2000 (replacement unit):	53,752.03
2001 (replacement unit):	49,705.96
2002 (replacement unit):	45,964.45
2003 (replacement unit):	42,504.58
	250,054.47 K$
Extension period - 1999 to 2005:	
1999 (1st replacement):	58,127.44
2000 (1st replacement):	53,752.03
2001 (1st replacement):	49,705.96
2002 (1st replacement):	45,964.45
2003 (1st replacement):	42,504.58

2004 (2nd replacement):	47,820.72
2005 (2nd replacement):	44,221.12
	342,096.31 K$

Nature of Financial Constraints

In the context of planning and evaluation, the term “financial constraints” is used for limits on the amount of funds an electric utility can obtain to support its capacity expansion plans. Broadly speaking, there are two types of financial constraints: profitability and liquidity.

Profitability

Profitability is the principal factor affecting the ability of a corporation to attract funds for capital investment. While often discussed in terms of historical performance, profitability in the present context is a prospective or forward-looking concept. It refers to the value of cash flows an investment will generate in relation to the required capital outlays. The future cash flows are those that well-informed investors can anticipate. In general, if investors can anticipate that an electric utility will make profitable investments, then that utility should be able to attract the funds necessary for capacity expansion.

Profitability depends on market as well as cost factors. Customers must be willing to pay for the electric services offered. In the electric utility industry, profitability also depends on regulation. Investment will be profitable if, among other things, investors anticipate that utility managers will make investments regarded by regulators as “prudent”, that regulators will permit the recovery of all “prudent” investment costs, and that earned rates of return will equal or exceed the cost of capital on average. Because profitability is forward-looking, explicit measures of profitability would require models of demand and regulation to solve for revenues and costs simultaneously.

In practice, profitability is often an assumption of capacity planning models, not a simulation result. The cost minimization objective assumes that utility rates are designed for full cost recovery, and that projected load growth reflects these cost-recovery rates. Profit incentives to follow other than least-cost expansion paths are not included. Accordingly, measures of profitability are seldom considered as candidate indicators of financial constraints. However, a measure of system-average rates will serve as a check on the reasonableness of profitability and load growth forecasts in light of regulatory and market responses to conventional cost recovery paths.

Liquidity

A company may be profitable in the sense used above and yet still confront financial constraints on capacity expansion. Liquidity constraints arise from the need to meet fixed cash obligations, such as those associated with interest and principal payments on debt financing. Failure to meet cash obligations, or even the prospect of failure, can disrupt normal business operations. The higher the proportion of capital raised through debt financing, the greater the likelihood that the utility will fail to make payments. Thus, one liquidity constraint is that debt financing does not exceed a given percentage of total capital requirements.

Some of the constraints imposed by debt financing are rigid. Debt contracts often restrict the amounts of additional debt that can be issued, specify minimum working capital requirements

and threshold interest coverage ratios, and limit the payments of cash dividends. Violation of these terms is a basis for technical default, although it may not trigger a full-scale response from creditors.

Other than the covenants contained in debt contracts, financial "constraints" are rarely hard and fast. More often, financial constraints are in the nature of warning signals that an investment program may require the utility to adopt an unconventional financial policy, one that is beyond the norms of the public utility industry. Although a policy "fix" of this sort may be possible, it might entail institutional costs that are perceived as unacceptable by top management. A premium is placed on adherence to the financial policies to which the securities markets have become accustomed. These norms cannot be rigorously specified. Thus, in most cases financial constraints are not strict numerical bounds; they should be interpreted as guidelines rather than rigid boundaries.

Price Constraints

One limitation of capacity planning models is that growth is typically viewed as being independent of the capacity expansion plan. In fact, however, load growth is influenced by the rates a utility charges, which in turn depend on the choice of a capacity expansion plan.

Thus, the conditions under which a utility violates apparent financial constraints may also presage or be coincident with two potential problems: (1) reductions in load caused by rate increases associated with additions of capital-intensive resources to utility rate bases; and (2) a closely linked phenomenon, that these same capital-intensive resources carry higher risks of partial disallowance from rate base. What are interpreted. as financial constraints may in fact be market or regulatory constraints.

Financial simulation models routinely assume that the utility wishes to follow conventional financing and ratemaking policies. This is important because a financial indicator value is based on the result of certain financing and ratemaking policies, and it is possible that alternative, nontraditional financing or ratemaking policies could be used to make apparent constraints more tolerable. To reiterate, financial constraints do not necessarily mean that a capacity expansion program is infeasible, but that special action may be required to make it feasible.

Candidate indicators of financial constraints can be classified in one of two categories: those that rely on information from the balance sheet and those that rely on information from the income statement and/or changes in financial position. This classification is helpful primarily because it is easier to compile the information required. for balance sheet ratios than it is to compile the information for income statement ratios. Indicators derived from the balance sheet include the earning assets ratio and the corporate debt-equity ratio, while indicators derived from income statements include the interest coverage ratio and the percentage of funds generated internally.

Feasibility Measures

EGEAS uses three feasibility measures: the earning assets ratio, the interest coverage ratio, and the percentage increase in system average cost. To calculate these measures, EGEAS needs information about all of a utility's assets, not just those associated. with new resources. Among the assets not normally modeled in EGEAS are existing resources, existing and future transmission and distribution facilities, and general plant assets that may be indirectly associated with generation expansion. These assets are lumped together and called "non-EGEAS" assets. For these non-EGEAS assets, you must supply annual values for depreciation, rate base and

CWIP balances, operating costs, and other electric revenues. In addition to information about non-EGEAS assets, you can also specify annual values for working capital to be allowed in rate base and the percentage of energy lost through transmission and distribution; the latter parameter is used to determine system sales from system demand on the resources.

Non-EGEAS Assets

To determine the revenue requirements associated with building and operating new resources, EGEAS needs only the capital costs of new resources and the operating costs of all resources. To calculate the economic feasibility of an expansion plan, however, EGEAS needs information about all of the utility's assets. In particular, EGEAS must account for costs associated with transmission and distribution assets and with existing resources which have already entered the rate base. These assets are classified as "non-EGEAS" assets.

Define:

y	=	year of study period
ZRETRB(y)	=	system rate of return on rate base (percent)
ZRPROP(y)	=	property tax rate (percent)
ZXDEPR(y)	=	non-EGEAS depreciation (K$)
ZXCWIP(y)	=	year-end non-EGEAS CWIP balance (K$)
ZXNERB(y)	=	year-end non-EGEAS rate base (K$)
ZXOPER(y)	=	non-EGEAS operating cost (K$)
ZXREVS(y)	=	other electric revenues (K$)
ZZWCAP(y)	=	working capital (K$)
ZYCWEX(y)	=	year-end non-EGEAS CWIP balance ($)
ZYRBEX(y)	=	year-end non-EGEAS rate base ($)
ZYCWMX(y)	=	mid-year non-EGEAS CWIP balance ($)
ZYRBMX(y)	=	mid-year non-EGEAS rate base
ZYFIXX(y)	=	non-EGEAS fixed charges ($)
ZYOPRX(y)	=	non-EGEAS operating cost ($)
ZYREVX(y)	=	other electric revenues ($)
ZYRBMW(y)	=	working capital ($)
Note: K$	=	thousands of current year dollars.

Save year-end values of non-EGEAS CWIP and rate base ($)

ZYCWEX(y)	=	1,000 * ZXCWIP(y)
ZYRBEX(y)	=	1,000 * ZXNERB(y)

Calculate mid-year values of non-EGEAS CWIP and rate base ($)

$$ZYCWMX(y) = \frac{ZYCWEX(y-1) + ZYCWEX(y)}{2}$$

$$ZYRBMX(y) = \frac{ZYRBEX(y-1) + ZYRBEX(y)}{2}$$

Calculate non-EGEAS rate base fixed charges ($)

$$ZYFIXX(y) = 1{,}000 * ZXDEPR(y) + ZYRBMX(y) * \left(\frac{ZRETRB(y)}{100} + \frac{ZRPROP(y)}{100} \right)$$

Save non-EGEAS operating cost ($)

ZYOPRX(y) = 1,000 * ZXOPER(y)

Save other electric revenues ($)

ZYREVX(y) = 1,000 * ZXREVS(y)

Save working capital ($)

ZYRBMW(y) = 1,000 * ZZWCAP(y)

Total Rate Base and CWIP

To calculate one or more economic feasibility measures for an expansion plan, EGEAS needs the mid-year values for the total system rate base and CWIP. EGEAS also needs total fixed and operating costs. These totals are obtained by adding the non-EGEAS and working capital values to the summation taken over all new resources.

Define:

y	=	year of study period
UKCWDR(y)	=	year-end CWIP balance for a resource - direct cost in rate base ($)
UKCWDN(y)	=	year-end CWIP balance for a resource - direct cost not in rate base ($)
UKCWAE(y)	=	year-end CWIP balance for a resource - equity AFUDC ($)
UKCWAD(y)	=	year-end CWIP balance for a resource - debt AFUDC ($)
UKZRBV(y)	=	year-end rate base value for a resource ($)
UKZFIX(y)	=	capital fixed charges for a resource ($)
ZYCWEE(y)	=	year-end CWIP balance - EGEAS assets ($)
ZYCWME(y)	=	mid-year CWIP balance - EGEAS assets ($)
ZYCWMX(y)	=	mid-year CWIP balance - non-EGEAS assets ($)
ZYCWMT(y)	=	mid-year CWIP balance - total ($)
ZYRBEE(y)	=	year-end rate base - EGEAS assets ($)
ZYRBME(y)	=	mid-year rate base - EGEAS assets ($)
ZYRBMX(y)	=	mid-year rate base - non-EGEAS assets ($)

ZYRBMW(y) = mid-year rate base - working capital ($)

ZYRBMT(y) = mid-year rate base - total ($)

ZYFIXE(y) = fixed costs - EGEAS assets ($)

ZYFIXX(y) = fixed costs - non-EGEAS assets ($)

ZYFIXT(y) = fixed costs - total ($)

ZYOPRE(y) = operating costs - EGEAS assets ($)

ZYOPRX(y) = operating costs - non-EGEAS assets ($)

ZYOPRT(y) = operating costs - total ($)

ZYTOTT(y) = total costs - total ($)

Determine CWIP balances ($).

Year-end CWIP balance for EGEAS assets. Values are summed over all resources under construction.

ZYCWEE(y) = sum of { UKCWAE(y) + UKCWAD(y) + UKCWDN(y) }

Mid-year CWIP balance for EGEAS assets.

$$ZYCWME(y) = \frac{ZYCWEE(y-1) + ZYCWEE(y)}{2}$$

Mid-year CWIP balance total

ZYCWMT(y) = ZYCWME(y) + ZYCWMX(y)

Determine rate base value ($):

Year-end rate base value for EGEAS assets. Values are summed over all resources under construction (if CWIP is allowed in rate base) and over all new resources that are operating.

ZYRBEE(y) = sum of [UKZRBV(y)]

Mid-year rate base value for EGEAS assets

$$ZYRBME(y) = \frac{ZYRBEE(y-1) + ZYRBEE(y)}{2}$$

Mid-year rate base value total

ZYRBMT(y) = ZYRBME(y) + ZYRBMX(y) + ZYRBMW(y)

Determine total fixed costs ($):

For EGEAS assets, values are summed over all resources.

ZYFIXE(y) = sum of [UKZFIX(y)]

Total fixed costs

ZYFIXT(y) = ZYFIXE(y) + ZYFIXX(y)

Determine total operating costs ($):

For EGEAS assets, operating cost is simply the annual production cost, including fixed operating and maintenance cost. Fuel and O&M costs are summed over all resources - existing, committed, and generic. Depending on options you specify, some detailed cost components may also be included.

ZYOPRE(y) = fuel cost + variable O&M cost
+ fixed O&M cost
+ purchased energy cost
+ unserved energy cost
+ emission allowance costs
+ detailed costs

Total operating costs

ZYOPRT(y) = ZYOPRE(y) + ZYOPRX(y)

Determine total costs - sum of fixed and operating costs ($):

ZYTOTT(y) = ZYFIXT(y) + ZYOPRT(y)

Earning Assets Ratio

The earning assets ratio reflects the fact that the assets supporting debt financing can be divided into two categories: those assets that are in rate base and hence generating cash returns – which may include some construction work in progress – and those assets that are not in rate base, for which there are only promises of future cash returns – mainly plants under construction. The ratio of earning assets to total assets is the balance-sheet equivalent of a ratio used by industry analysts, the ratio of net income excluding AFUDC to net income. The Earning Assets Ratio will decline as construction of a capital intensive resource progresses and then rise again as the resource goes into rate base. EGEAS uses mid-year values for rate base and CWIP in calculating the earning assets ratio; values for both new resources and the non-EGEAS assets are included.

The earning assets ratio is calculated as the ratio of earning assets (those in rate base) to total assets (rate base plus CWIP). Rate base and CWIP values for non-EGEAS assets and working capital are included as well as the totals for all new resources.

Define:

ZYRBMT(y) = mid-year rate base ($) in year y

ZYCWMT(y) = mid-year CWIP balance ($) in year y

EASRAT(y) = earning assets ratio in year y

Calculate earning assets ratio:

$$EASRAT(y) = \frac{ZYRBMT(y)}{ZYRBMT(y) + ZYCWMT(y)}$$

When calculating a multiple-year average for the earning assets ratio, EGEAS sums the earning assets (numerators) and sums the total assets (denominators) over all the included years and then

takes the ratio. This is not the same as taking the average of the individual year ratios. The definition used by EGEAS is dictated by the convexity condition required for convergence with a Benders' decomposition optimization.

Interest Coverage Ratio

Liquidity constraints can be expressed in terms of accruals or cash flows – information from the income or funds flow statements – rather than asset and liability values. A very common measure of liquidity is the interest coverage ratio, which relates earnings before interest and taxes to interest payments on fixed-income obligations. It is an indicator of the margin for error on debt obligations. In a sense it is a more immediate indicator of financial constraints than the balance sheet items, because it focuses on current values only. In EGEAS, the interest coverage ratio is expressed in terms of accruals rather than cash items only.

Interest coverage is calculated as the ratio of the earnings before interest and taxes (EBIT) to the debt expense. The numerator of the ratio (EBIT) is determined by applying the system rate of return on rate base to the total mid-year value of rate base. Included in the total rate base are the non-EGEAS rate base and working capital values, as well as the sum of the rate base values for all the new resources. The denominator of the ratio, the debt expense, is obtained by applying the debt interest rate to the debt balance. The debt balance is determined by applying the percent of debt in the capital structure to the sum of mid-year rate base and CWIP values.

Define:

ZRETRB(y) = system rate of return on rate base (percent) in year y

ZRDEBT(y) = percent debt in capital structure in year y

ZRDINT(y) = debt interest rate (percent) in year y

ZYRBMT(y) = mid-year rate base ($) in year y

ZYCWMT(y) = mid-year CWIP balance ($) in year y

COVRAT(y) = interest coverage ratio in year y

Calculate earnings before interest and taxes ($):

ZYEBIT(y) = ZRETRB(y) * ZYRBMT(y)

Calculate debt balance ($):

$$ZYDBAL(y) = \frac{ZRDEBT(y)}{100} * (ZYRBMT(y) + ZYCWMT(y))$$

Calculate debt expense ($):

$$ZYDEXP(y) = \frac{ZRDINT(y)}{100} * ZYDBAL(y)$$

Calculate interest coverage ratio:

$$COVRAT(y) = \frac{ZYEBIT(y)}{ZYDEXP(y)}$$

When calculating a multiple-year average for interest coverage, EGEAS sums the EBIT's (numerators) and sums the interest expenses (denominators) over all the included years and then

takes the ratio. This is not the same as taking the average of the individual year ratios. For example, if N_i represents the EBIT and D_i represents the debt expense for year i, then the interest coverage ratios for years 1 and 2 would be:

$$R_1 = \frac{N_1}{D_1} \qquad R_2 = \frac{N_2}{D_2}$$

and the 2-year average calculated by EGEAS would be:

$$R_{12} = \frac{N_1 + N_2}{D_1 + D_2}$$

Note that R_{12} is not the same as:

$$A = \frac{R_1 + R_2}{2}$$

In all cases, R_{12} and A both have values between R_1 and R_2 and in most cases, their values are fairly close. EGEAS uses R_{12} rather than A so that multiple-year interest coverage and earnings assets ratio constraints can be applied to the Benders' decomposition optimization. If A were used, then the optimization might not converge to the correct solution because the convexity condition would be violated.

System Average Cost

Because rate spikes are often associated with additions of large resources to utility rate bases, a load reduction may coincide with plant additions. Thus, another candidate indicator that is useful to follow – a market or rate indicator, not strictly a financial indicator – is the annual increase in system average cost based on conventional ratemaking procedures.

Calculation of System Average Rate

System average rate can be calculated with any of the cost analysis format options: 1 - levelized fixed charges, 2 - annual fixed charges, and 3 - construction costs. To calculate the system average cost, EGEAS sums the capital fixed charges (calculated in accordance with the particular cost analysis format) and operating costs (including the costs associated with non-EGEAS assets), subtracts other electric revenues (such as wheeling revenues), and divides by the system electricity sales. System sales can be determined by subtracting the energy provided by demand side options from the system demand and then accounting for line losses. In addition to calculating the system average rate, EGEAS also determines the percentage increase in rate from year to year and the compound average increase over several years.

Define:

ENERGY(y)	=	annual system load (GWh) in year y
EDSM(y)	=	energy contributed by demand-side management resource (GWh) in year y
ZZLLOS(y)	=	line loss (percent of demand) in year y
ZYTOTT(y)	=	total fixed plus operating cost ($) in year y
ZYREVX(y)	=	other electric revenues ($) in year y

AVGRAT(y) = system average rate ($/MWh or mills/kWh) in year y

ARINCR(y) = single year increase in system average rate, from year y1 to year y (percent)

ARINCM(y) = compound average increase in system average rate over N years

Calculate system sales (GWh):

$$ZYSSAL(y) = (ENERGY(y) - EDSM(y)) * \frac{100 - ZZLLOS(y)}{100}$$

Calculate system average rate ($/MWh):

$$AVGRAT(y) = \frac{ZYTOTT(y) - ZYREVX(y)}{ZYSSAL(y)}$$

Single year increase (percent):

$$RAT = \frac{AVGRAT(y)}{AVGRAT(y - 1)}$$

$$ARINCR(y) = 100 * (RAT - 1)$$

Compound average increase over N years (percent):

$$RATM = \left(\frac{AVGRAT(y)}{AVGRAT(y - N)} \right)^{1/N}$$

$$ARINCM(y) = 100 * (RATM - 1)$$

Calculation of Levelized System Average Rate

With the dynamic program, you can specify an option of using the levelized system average rate, instead of the present worth of revenue requirements, as the objective function for expansion optimization. The levelized system average rate is defined as the present worth sum of the system average rate in each year of the study period. The EGEAS program calculates the system average rate in each year, discounts it back to the base year, sums over all study period years, and then converts the sum into a levelized value.

Define:

y = study period year
(1 for 1st study year, 2 for 2nd year, etc.)

Y1 = first study year

IDBSYR = base year of data base

DISC = system discount rate (percent)

AVGRAT(y) = system average rate for year y

RATPW(y) = present worth of system average rate for year y

Discount the annual system average rate to the year immediately preceding the first study period year:

DR = DISC/100

RATPW(y) = AVGRAT(y)/(1.0 + DR)^{y}

Note that if you select the levelized system rate extension option, then the extension period is included in the summation and levelization process. For the extension period, EGEAS calculates the system average rate in base year dollars, so the extension period rate must be escalated from the base year to the year before the first study period year.

M = (Y_1 - 1) - IDBSYR

RATPWEXT = AVGRATEXT * (1.0 + DR)^{M}

Calculate the levelized system rate:

N = number of years included in levelization (including extension period if requested)

SARLEV = [summation of RATPW over N years] $* DR * \frac{(1+DR)^{N}}{(1+DR)^{N}-1}$

If the extension period is infinite, the formula becomes:

SARLEV = [summation of RATPW] * DR

Customer Revenue Accounting

Revenue calculations may be of importance in a typical IRP analysis involving DSM programs. Customer class rates are allowed so that the total revenues from the system load profile can be derived. If the DSM impacts total revenues, then you can evaluate these changes within EGEAS by including the DSM program in the optimization process.

EGEAS allows you to supply hourly load data for individual class load components for System A; possible classes include residential, commercial, industrial, street light, and other. You have complete freedom in subdividing the load. EGEAS combines the component loads to obtain the total system A load using probabilistic load modification in the same manner as it constructs the pool load when interconnections are modeled.

Total revenues may include revenues derived from demand charges, energy charges, and customer charges. After you have supplied multiple load components, you can then specify a demand, energy or customer charge rate with a trajectory for each rate. For every study period year (y), EGEAS calculates the revenues based on the class load components and the rates specified for each class component (i).

Define:

RYDPEK(i) = annual peak load (MW)

RYDDSM(i) = DSM capacity (MW)

RVDLOS(i) = line losses - percent of demand (percent)

RYDRAT(i) = demand charge ($/kW/month), escalated from base year to year y

RYENRG(i) = annual energy (GWh)

RYEDSM(i) = DSM energy (GWh)

RVELOS(i) = line losses - percent of energy (percent)

RYERAT(i) = energy charge (cents/kWh), escalated from base year to year y

NCUSRY(i) = number of customers, escalated from base year to year y

RYCRAT(i) = customer charge ($/month/customer), escalated from base year to year y

The following calculations are performed in EGEAS to obtain the revenues.

Demand revenues:

demand losses (MW):

$$RYDLOS = (RYDPEK - RYDDSM) * \frac{RVDLOS}{100}$$

net demand (MW):

DNET = RYDPEK - RYDDSM - RYDLOS

demand revenues ($):

RYDREV = RYDRAT * DNET * 12,000

Note: 12,000 = 12 (month/year) * 1,000 (kW/MW)

Energy revenues:

energy losses (GWh):

$$RYELOS = (RYENRG - RYEDSM) * \frac{RVELOS}{100}$$

net energy (GWh):

ENET = RYENRG - RYEDSM - RYELOS

energy revenues ($):

RYEREV = RYERAT * ENET * 10,000

Note: 10,000 = 0.01 ($/cents) * 1,000,000 (kWh/GWh)

Customer revenues ($):

RYCREV = RYCRAT * NCUSRY * 12

Note: 12 = 12 (month/year)

Total revenues ($):

RYTREV = RYDREV + RYEREV + RYCREV

8
SYSTEM INTERCONNECTIONS

Overview

Most electric power systems are involved to some extent in reserve sharing and economy interchange with neighboring systems. The importance attached to the effects of interconnections on expansion planning of individual power systems varies widely among utilities. The EGEAS interconnection model is designed to accommodate this wide variety of utility requirements by providing several levels of detail and sophistication to meet different utility needs and data availability. Up to ten different utility systems can be represented in EGEAS (Systems A to J).

The EGEAS interconnection model is compatible with both the dynamic programming and pre-specified pathway analysis options. In particular, the methodology is based on probabilistic production costing. Reserve requirements, as affected by interconnections, are treated separately from economy interchange. You can specify either or both of these options. EGEAS can model economy interchange at a subperiod level of detail; maintenance scheduling is available for Systems B-J as well as for System A.

In all cases, the primary electric power system (System A) is optimized and the characteristics of the neighboring systems (System B-J) are assumed to be fixed. For the most detailed formulations of the interconnection model, Systems B to J generation expansion plans must be specified in detail. Capacity limits for the tie-lines (transmission lines connecting the systems) must also be specified.

Reserve Sharing

When two or more systems are interconnected, the reserves are normally pooled so that emergency assistance is available to one system up to the limit of generating capacity reserve available in the remaining system. Such reserve sharing generally increases system reliability. The amount of reserves available from the pool may be limited either by transmission tie capacity or by the probability that reserve capacity from the pool is available when needed.

The reserve sharing facet of the EGEAS interconnection model is used in determining reserve margins and loss of load probabilities. It has no effect on production cost or unserved energy calculations. For the reliability calculations, the contribution of Systems B to J is treated as a single generating unit. Three methods are provided for determining the characteristics of this unit:

Method 1 – Tie-Line Capacity Limit

A simple model is justified for a class of small systems connected to much larger utilities or power pools. In these cases, reserves are available with almost perfect reliability up to the limits of the tie-line capacity. For these systems, Systems B to J can be modeled as a generating unit with no forced outage and a rating equal to the tie-line capacity. For this simplification to be valid, the tie-line capacity must be small relative to the available reserves. The advantage of this option is that it requires no data for Systems B to J.

Method 2 – Percent Load Limit

If a small system is connected to a much larger system with a large tie-line capacity, the neighboring system can sometimes supply the entire load. In reality, however, the utility would not want to depend solely on reserves from outside its system. To overcome this problem, EGEAS limits the interconnected load to a user-specified percentage of the System A peak load. Thus, the contribution of Systems B to J is modeled as a generating unit with no forced outage and a rating equal to a given percentage of the peak load.

Method 3 – Multi-State Unit

For this option, you must develop a probability distribution for available capacity on the other systems. Subtracting the Systems B-J peak loads from this distribution gives the probability distribution for available reserves. This resulting distribution is then represented as a generating unit with multiple outage states. If necessary, the capacity of the unit is truncated to the value of the tie-line capacity. You specify the characteristics of this unit based on conditions at the time of the Systems B to J peak loads. More reserves may be available at other times in the year, but reliability constraints are most sensitive to peak load conditions, and Systems B to J maintenance would tend to levelize reserves during off-peak periods.

Following is an example of a multi-state unit:

Capacity, MW	Probability
0	0.10
100	0.22
200	0.30
300	0.20
400	0.12
500	0.06

If the tie-line limit were 700 MW, then 6 percent of the time the reserve would be 500 MW, 12 percent of the time, it would be 400 MW, and so forth; 10 percent of the time, there would be no reserves available. If the tie-line limit were 350 MW, then 18 percent of the time the reserve would be 350 MW (truncated from 500 MW and 400 MW), 20 percent of the time, it would be 300 MW, and so forth.

Economy Interchange

In the EGEAS interconnection model, the system dispatcher's hour-by-hour decisions are replaced by simulated interchanges based on a probabilistic determination of the availability and cost of economy interchanges. The basic methodology involves multiple production cost dispatches: one for each of the individual systems (independent dispatches) and one for the combined systems (pool dispatch). If there are tie-line limits, EGEAS converts tie-line capacity limits into energy limits and imposes an energy constraint. EGEAS then performs cost accounting to determine costs of the transferred energy.

Methodology

The basic methodology for economy interchange involves multiple production cost dispatches: one for each of the individual systems (A to J) supplying the load within each system, and one for the combined systems supplying the combined load from all systems. The combined load accounts for correlations between loads on the individual systems. The total saving in production cost (fuel, variable O&M, and unserved energy costs) is the difference between the sum of the individual dispatch costs and the combined system dispatch cost.

EGEAS measures the amount of energy transferred between the systems by looking at changes in the amount of energy generated by individual resources when switching from the independent to the pool dispatch. If a resource generates more energy in the pool dispatch, the increase is treated as a sale (the extra energy is sold to the other system to meet its load). If a resource generates less energy, the decrease is treated as a purchase (the lost generation must be replaced by energy purchased from the other system). By summing the increases and decreases over all the resources on System A, EGEAS obtains the amount of energy sold and purchased by System A. Similarly, EGEAS determines the energy sold and purchased by Systems B to J. The sales and purchases computed by EGEAS may be less than the true amounts determined by an hourly dispatch model because the sales and purchases during a subperiod may offset each other in some units. However, knowledge of the true purchase and sale amounts is not necessary for determining the savings from economy interchange unless tie-line limits are binding.

To illustrate how EGEAS calculates purchase and sales, consider two systems with the following characteristics:

System A	-	Resources A1, A2, A3
System B	-	Resources B1, B2, B3
Merit loading order	-	A1, B1, A2, B2, B3, A3
For all resources: Capacity	=	100 MW
Forced outage rate	=	0

For a particular hour, assume the load on System A is 70 MW and the load on System B is 280 MW. Then the generation in the independent and pool dispatches is:

Hour 1	**Generation, MW**					
	A1	**A2**	**A3**	**B1**	**B2**	**B3**
Independent	70	0	0	100	100	80
Pool	*100*	*100*	*0*	*100*	*50*	*0*
Difference	+30	+100	0	0	-50	-80

At this hour, System A is selling 30 + 100 = 130 MWH and System B is purchasing 50 + 80 = 130 MWH.

For another hour, assume the load on System A is 270 MW and the load on System B is 160 MW. Then the generation in the independent and pool dispatches is:

Hour 2	Generation, MW					
	A1	**A2**	**A3**	**B1**	**B2**	**B3**
Independent	100	100	70	100	60	0
Pool	*100*	*100*	*0*	*100*	*100*	*30*
Difference	0	0	-70	0	+40	+30

At this second hour, System A is purchasing 70 MWH and System B is selling 40 + 30 = 70 MWH.

If the two hours are combined, as is the case when EGEAS uses load duration curves, the generation values are:

Hours 1 & 2	Generation, MW					
	A1	**A2**	**A3**	**B1**	**B2**	**B3**
Independent	170	100	70	200	160	80
Pool	200	200	0	200	150	30
Difference	+30	+100	-70	0	-10	-50

Looking just at the combined results, EGEAS would compute the System A sales as 30 + 100 = 130 MWH and the System A purchases as 70 MWH. But for System B, EGEAS would find no sales and would determine the purchases to be 10 + 50 = 60 MWH. EGEAS finds no sales for System B because the increases for Resources B2 and B3 in Hour 2 are offset by greater decreases in Hour 1. Note that the net purchases (purchases less sales) for System A (+60) and for System B (-60) are in agreement.

Tie-Line Limits

Because EGEAS does not perform production costing on an hourly basis, it cannot impose tie-line constraints directly. Instead, EGEAS converts tie-line capacity limits into energy limits and imposes an energy constraint. The formula for the energy limit is:

Energy Limit = Capacity * (Capacity Factor/100) * Hours

where the capacity and capacity factor are input. Separate limits are computed for sales and purchases. You should choose capacity factors that minimize the chances of violating the capacity constraint while still allowing as much economic transfer as possible. Nevertheless, the energy constraint does not guarantee that the capacity limit is satisfied. Thus, the modeling of tie-line capacity limits is only approximate because EGEAS does not use an hourly model and because EGEAS may miss purchases and sales that partially offset one another.

If a tie-line purchase energy or sale energy constraint is violated by the pool dispatch, then EGEAS adjusts the generation of individual resources to reduce the transfers to the allowed levels. If the sales from System A to Systems B to J are too great, then EGEAS reduces the generation by System A resources and increases the generation by Systems B to J resources. Similarly, if the purchases by System A from Systems B to J are too great, then EGEAS increases the generation by System A and reduces the generation by Systems B to J. In all cases,

EGEAS works only with resources whose generation is different in the independent and pool dispatches; the final adjusted energy is always between the values from the independent and unadjusted pool dispatches. When reducing resource generation, EGEAS always starts with resources having the highest marginal operating cost. When increasing generation, EGEAS starts with resources having the lowest cost. Thus, EGEAS retains as much of the economy interchange savings as possible.

To illustrate the application of tie-line energy constraints, consider the example given above. Assume that the sales (of System A) are limited by the following parameters:

Capacity limit = 200 MW

Capacity factor limit = 30 percent

For the two hours shown above the energy limit is:

Energy limit = (200) * (30/100) * 2 = 120 MWH

The sales for System A (130 MWH) exceed the energy limit by 10 MWH, so the sales must be reduced. The most expensive System A resource with sales is Resource A2, so its generation is reduced from 200 MWH to 190 MWH. For System B, the calculated purchases are only 60 MWH, but because the sales for System A were reduced, the System B purchases must also be decreased. The least expensive System B resource with purchases is Resource B2, so its generation is raised from 150 MWH to 160 MWH. The resulting values are:

Hours 1 & 2	**Generation, MW**					
	A1	**A2**	**A3**	**B1**	**B2**	**B3**
Independent	170	100	70	200	160	80
Pool	*200*	*200*	*0*	*200*	*150*	*30*
Difference	+30	+100	-70	0	-10	-50

Cost Accounting

After the multiple production costing dispatches have been simulated and after any adjustments in the pool dispatch for tie-line limits are made, EGEAS performs cost accounting. EGEAS provides four different pricing options for assigning costs to the transferred energy. The four options are: split savings, cost plus profit percentage, billing rates and buy and sell rates. Each of the four options is discussed below.

Split savings: The overall savings, measured as the difference in cost between dispatching all the systems combined as a pool and dispatching the systems independently, are split equally among all the systems. If there are two systems, the sale price is midway between two running rates so the total net savings are split equally between the two systems; the seller obtains more than its cost from the sale and the buyer benefits by replacing expensive generation with less expensive purchased energy. The following table illustrates the split savings calculations for two systems.

	System A	System B	Total
MWH Load	10,000	8,000	18,000
Independent Dispatch			
Generation, MWH	10,000	8,000	18,000
Cost, $	100,000	72,000	172,000
Pool Dispatch			
Generation, MWH	7,000	11,000	18,000
Cost, $	70,000	99,000	169,000
Savings	1,500	1,500	3,000
Net Cost	98,500	70,500	169,000

In this example, the resources on System B are less expensive than those on System A, so in the pool dispatch, System B sells 3,000 MWH to System A. The overall savings are $3,000 ($172,000 - $169,000). These savings are split equally between the systems, so the cost for System A is its independent dispatch cost ($100,000) less the savings ($1,500) or $98,500. For System B, the cost is $72,000 - $1,500 = $70,500. In terms of where the money is spent, the cost for System A is:

Generation cost	$70,000
Purchases	28,500
Net cost	$98,500

For System B, the net cost is:

Generation cost	$ 99,000
Purchases	(28,500)
Net cost	$70,500

Costs plus profit percentage: With this option, EGEAS firsts determines the marginal operating cost for all resources with increased generation in the pool dispatch, then adds the profit percentage specified, and then uses the resulting value as the price of the energy sold. The purchasing system must pay the selling system at a rate based on a specified percentage above the operating cost of the resources supplying the energy. If the profit percentage is too high, it may no longer be cost effective for one of the systems to participate in the interchange. If this situation is detected during the analysis, EGEAS will provide a warning message.

Billing rates method: With this option, each system is classified as a net seller or a net purchaser of energy. For each net seller, EGEAS calculates a sell rate, which is the average cost of the energy sold. It also determines the average sell rate over all net sellers. For each net purchaser,

EGEAS calculates a buy rate, which is the average cost of the displaced generation. It also determines the average buy rate over all net purchasers. EGEAS then calculates a billing rate for each net seller as the average of that system's own sell rate and the average buy rate from all the net purchasers. The energy sold by this system will be priced at this rate. Similarly, EGEAS calculates a billing rate for each net purchaser as the average of that system's own buy rate and the average sell rate over all the net sellers. The energy purchased by this system will be priced at this rate.

Define:

GEI(i)	=	generation (GWh) for area i, independent dispatch
UEI(i)	=	unserved energy (GWh) for area i, independent dispatch
ZIND(i)	=	(fuel + variable O&M + unserved energy) costs (K$) for area i, independent dispatch
UEITOT	=	total unserved energy (GWh) summed over all areas, independent dispatch
UEITOT	=	_ UEI (i)
GEP(i)	=	generation (GWh) for area i, pool dispatch
UEPTOT	=	total unserved energy (GWh) for pool dispatch
UEP(i)	=	unserved energy (GWh), for area i, pool dispatch
UEP(i)	=	UEPTOT * [UEI(i) /UEITOT]
ZPOL(i)	=	(fuel + variable O&M + unserved energy) costs (K$) for area i, pool dispatch

Net sales for each area and differences in cost:

DGE(i)	=	GEP(i)	-	GEI(i)
DUE(i)	=	UEP(i)	-	UEI(i)
DE(i)	=	DGE(i)	+	DUE(i)
DZ(i)	=	ZPOL(i)	-	ZIND(i)

Sale and buy rates, for each area:

RAT(i)	=	DZ(i)	/	DE(i)

Average sell rate:

DES	=	_ DE (i)
DZS	=	_ DZ (i) for all areas that are net sellers: DE(i) > 0
RATS	=	DZS / DES

Average buy rate:

DEP	=	- _ DE (i)
DZP	=	- _ DZ (i) for all areas that are net purchasers: DE(i) < 0

RATP = DZP / DEP

Savings for a net selling area, DE(i) > 0

$$FUSAV(i) = \left(\frac{RAT(i) + RATP}{2} * DE(i) \right) - DZ(i)$$

Savings for a net purchasing area, DE(i) < 0

$$FUSAV(i) = \left(\frac{RAT(i) + RATS}{2} * DE(i) \right) - DZ(i)$$

Net production cost for pool dispatch, by area:

ZNETSS(i) = ZIND(i) - FUSAV(i)

Net purchased energy cost:

ZSPRCH(i) = ZNETSS(i) - ZPOL(i)

The following example illustrates the calculations using the billing rates method.

System	Generation (GWh)			Unserved Energy (GWh		
	Indep.	Pool	Diff.	Indep.	Pool	Diff.
	(GEI)	(GEP)	(DGE)	(UEI)	(UEP)	(DUE)
A	28868	28279	-589	419	6	-413
B	15606	15059	-547	11	0	-11
C	8240	10153	+1913	122	1	-121
D	4007	4850	+843	62	1	-61
E	9480	8521	-959	55	0	-55

	Cost (K$)		
System	Indep.	Pool	Difference
	(ZIND)	(ZPOL)	(DZ)
A	597,000	529,000	-68,000
B	435,000	405,000	-30,000
C	130,000	164,000	+34,000
D	71,000	79,000	+8,000
E	279,000	233,000	-46,000

Net sales, by area:

System	Net Sales DE (GWh)	Cost Difference DZ (K$)	Sell Rate ($/GWh)	Purchase Rate ($/GWh)
A: -589 – 413 =	-1,002	-68,000		67,864
B: -547 – 11 =	-558	-30,000		53,763
C: 1,913 – 121 =	+1,792	+34,000	18,984	
D: 843 – 61 =	+782	+8,000	10,230	
E: -959 – 55 =	-1,014	-46,000		45,365

Average sell rate:

DES = 1,792 + 782 = 2,574 GWh

DZS = 34,000 + 8,000 = 42,000 K$

RATS = 42,000 * 1,000 / 2,574 = 16,317 $/GWh

Average buy rate:

DEP = 1,002 + 558 + 1,014 = 2,574 GWh

DZP = 68,000 + 30,000 + 46,000 = 144,000 K$

RATP = 144,000 * 1,000 / 2,574 = 55,944 $/GWh

Savings, net production cost, net purchased energy cost, by area:

System A:

FUSAV(A) = 0.5 * [67,864 + 16,317] * (-1,002)/1,000

- (-68,000) = 25,825 K$

ZNETSS(A) = 597,000 - 25,825 = 571,175 K$

ZSPRCH(A) = 571,175 - 529,000 = 42,175 K$

System B:

FUSAV(B) = 0.5 * [53,763 + 16,317] * (-558)/1,000

- (- 30,000) = 10,448 K$

ZNETSS(B) = 435,000 - 10,448 = 424,552 K$

ZSPRCH(B) = 424,552 - 405,000 = 19,552 K$

System C:

FUSAV(C) = 0.5 * [18,984 + 55,944] * (1,792)/1,000

- (34,000) = 33,135 K$

ZNETSS(C) = 130,000 - 33,135 = 96,865 K$

ZSPRCH(C) = 96,865 - 164,000 = -67,175 K$

System D:

FUSAV(D) = 0.5 * [10,230 + 55,944] * (782)/1,000
- (8,000) = 17,874 K$

ZNETSS(D) = 71,000 - 17,874 = 53,126 K$

ZSPRCH(D) = 53,126 - 79,000 = -25,874 K$

System E:

FUSAV(E) = 0.5 * [45,365 + 16,317] * (- 1,014)/1,000
- (-46,000) = 14,727 K$

ZNETSS(E) = 279,000 - 14,727 = 264,273 K$

ZSPRCH(E) = 264,273 - 233,000 = 31,273 K$

Buy and sell Rates method: With this option, EGEAS calculates a sell rate and a buy rate for each system. The sell rate is the average cost of the energy sold and the buy rate is the average cost of the generation displaced by purchases. EGEAS also calculates system-wide sell and buy rates from the individual system values, weighted by the amount of energy sold or purchased. Next, EGEAS calculates, for each system, the cost of its purchases and the revenue from its sales. The purchase price is the average of the system's buy rate and the overall system sell rate; this price is applied to the amount of energy purchased. Similarly, the sale price is the average of the system's sell rate and the overall system buy rate; this price is applied to the amount of energy sold. This method is similar to the billing rates method, except that the sales and purchases for each system are treated separately instead of classifying the system as a net seller or net purchaser.

Input Requirements

The input requirements for the various options of the EGEAS interconnection model are given in Table 8-1. Items required for the analysis of a single stand-alone system are not included.

Table 8-1
Input Requirements for Interconnections

	Reserve Sharing			
Data Required	**Tie Limit**	**Percent Load**	**Multi-State Unit**	**Economy Interchange**
Tie-Line				
Purchase Limit	X		X	X
Tie-Lin Sale Limit				X
Percent Load Limit		X		
System B Reserve Capacity			X	
System B Generating Units				X
System B Load				X

9
ENVIRONMENTAL COMPLIANCE PLANNING

Overview

Environmental concerns can have a significant effect on how power plants are operated and also on which types of power plants should be built in the future. Federal and state regulators require utilities to consider the effect of power plants on air and water quality, ambient water temperature, and water supply. EGEAS can calculate several of these environmental effects when simulating an expansion plan. With the dynamic programming method, EGEAS can even impose environmental constraints so the resulting expansion plan is environmentally acceptable; environmental constraints are not allowed with the Bender's Decomposition option.

Because they are the primary source of SO_2 emissions and second only to transportation as a NO_X source, utilities must examine their resource plans to ensure compliance with the stringent emission limits enacted under the Clean Air Act Amendments (CAAA). Moreover, utilities must attain these limits within the overall integrated resource planning (IRP) framework. The impact of these efforts on utility operations can be profound.

Within the IRP process, resources must be evaluated with regard to their value in helping meet air quality standards. These evaluations include:

Short-term tactics that a utility can implement in a three- to five-year time frame, such as: demand-side management (DSM), changes in the dispatch order, installation of scrubbers, and fuel switching; and

Long-term strategies, such as plant retirements and acquisition of "clean" resources.

For a pre-specified pathway analysis, emissions may be reduced in one or more of the following ways:

Reduce the output of highly-polluting units by penalizing them with emission costs, and hence moving them later in the dispatch order.

Reduce the output of highly polluting units by placing upper limits on their allowed emissions.

Switch units to cleaner fuels.

Purchase more "clean" energy from another area.

For a dynamic programming optimization, emissions may be reduced by using any of the methods listed above for a pathway run, or by using one of the following methods:

Retire a highly-polluting unit.

Install a new "clean" unit; older less-efficient and more polluting units will end up generating less energy.

Add a DSM program to lower the system demand.

Refurbish an existing unit by adding pollution control equipment (e.g., a scrubber).

Switch a coal unit to burning low-sulfur coal.

Switch a coal unit to burning gas.

Reject plans that have unacceptably high emissions.

Assign a cost to emissions, thus economically eliminating plans with high emissions.

With the Benders' decomposition, EGEAS can calculate various environmental effects for an expansion plan, but cannot take any measures for reducing emissions.

Plant Emissions

EGEAS gives you considerable flexibility in modeling plant emissions. In particular, you define the type of emissions to be calculated and the units used for measuring the emissions. An emission type can be a specific chemical, such as sulfur dioxide (SO_2), a group of related compounds, such as nitrous oxides (NO_X), or the heat released to a neighboring body of water. An emission type can also represent a resource consumed by a plant, such as cooling water. Examples of emissions units include tons (of nitrous oxides), MBTU (of heat released), and millions of gallons (of water consumed). In all cases, you must supply information for determining the amount of emissions produced (or consumed) by the plant. EGEAS divides emissions into two classes: input and output.

Input class emissions represent the release to the environment of materials contained in the fuel consumed by the plant. Examples include sulfur dioxide (SO_2) and nitrous oxides (NO_X). For each input class emission, you must specify the amount in each fuel type (e.g., tons of SO_2 per ton of coal) and the emissions production rates for all resources that burn fuel containing that emission type. The production rate specifies the fraction of the material originally in the fuel that is released to the environment (e.g., tons of SO_2 released per ton of SO_2 in the fuel consumed). Thus, the production rate reflects the pollution control equipment installed for the resource. The formula for calculating input class emissions is:

Emissions (emissions mass) = Generation (GWh)

* Heat rate (MBTU/GWh)

/ Heat content of fuel (MBtu/fuel mass)

* Emission content of fuel (emissions mass/fuel mass)

* Emission production rate (fraction released)

As an example, consider a coal unit that has dry scrubbers which are 70 percent efficient and that burns 11,000 Btu/pound (or 22 MBtu/ton) coal with a sulfur content of one percent. SO_2 is produced at a rate of 2 pounds of SO_2 for each pound of sulfur in the coal. Thus, a one-percent sulfur content coal produces 40 pounds of SO_2 per ton or .02 tons of SO_2 per ton of coal. Because 30 percent of the SO_2 is released into the environment, the emissions production rate is 0.3. The emission produced by the unit amount to:

(0.02) * (0.3) = 0.006 ton SO_2/ton coal

If the unit has a heat rate of 8,000 MBtu/GWh and generates 1,000 GWh, the resulting emissions are:

(1,000 GWh) * (8,000 MBtu/GWh)/(22 MBtu/ton coal)

* (0.006 ton SO_2/ton coal)

= 1,056,000 ton SO_2

Output class emissions represent effects on the environment related to the generation output of the plant rather than amount of fuel burned. Examples include the amount of cooling water consumed or the heat output to a body of water. For output class emissions, the production rate specifies the quantity of emissions per GWH of generation (e.g., MBtu of heat output per GWH of generation). The formula for calculating output class emissions is:

Emissions (emissions mass) = Generation (GWh)

* Emission production rate (emissions mass/GWh)

Note that you can convert an input class emission type to output class by changing the emission production rate as follows:

Emission production rate (output class: emission mass / GWh) =

Heat rate of unit (MBtu/GWh)

/ Heat content of fuel (MBtu / fuel mass)

* Emission content of fuel (emission mass / fuel mass)

* Emission production rate (input class: fraction)

A particular emission type can be modeled as an input class emission for some generating units and as an output class emission for others. Note that there is a difference between the two classes in regards to multiple-block units. For an output class emission, the emissions per unit of generation (e.g., ton SO_2/GWh) is the same for all blocks, but for an input class emission, the emissions per unit of generation varies with the heat rate and hence can be different for each block.

Generic Sites

EGEAS allows you to define one or more generic sites for locating generating facilities. A generic site can represent a specific location, a type of location (such as a river suitable for wet cooling tower), or any arbitrary grouping of units. A site can also be used to represent Phase I or Phase II targeted units for emission compliance. You specify which units should be included at each site.

Emission Limits

For each defined emission type, you may specify emission limits at several different levels. First, you can specify emission limits for each generating unit. You must provide these limits in terms of mass units of emission per year. If you have the limits in other units (e.g., lb. SO_2/MBtu), you must convert the limits for input to EGEAS (e.g., tons SO_2/year). You may change the limits from year to year and you may allocate the limits to individual segments of the year.

Second, you can specify limits for each generic site. As with the unit limits, these limits are in terms of mass units of emission per year. There is less flexibility with site limits, however, because the site limits cannot be changed from year to year. Also, you cannot control how the limits are allocated to individual subyearly time periods. When EGEAS performs production costing at the subperiod level, the annual limits are always allocated to individual segments in proportion to the lengths of the segments. Thus, a segment with 4 weeks will be allocated 4/52 of the annual emissions limit.

Third, you can specify emission limits for the system as a whole. As with the individual unit limits, these limits can be changed from year to year and you can allocate the limits to different segments of the year. With multi-area modeling, you can specify system-wide limits only for the primary system (System A).

For all levels of detail, you can specify limits for some emission types and not for others. For the unit and system limits, you can also specify a starting year; before that year, the emissions are assumed to be unlimited. When no limit is specified for a particular emission type at a particular level of detail in some year, EGEAS still calculates the amount of emissions, but it cannot impose a constraint or calculate an emission cost associated with that quantity.

Allowance Costs and Credits

One of the ways EGEAS can reduce emissions is by assigning costs to emissions, thus providing an economic incentive to reduce generation by highly-polluting units. You can specify allowance costs and credits either for individual generating units or for the total system. These costs and credits are given in terms of dollars per unit of emissions (e.g., \$/ton SO_2). If you want to use costs but not credits, you can set the allowance credit to zero. Costs and credits can be applied after a production costing dispatch in one of three ways.

At the unit level, each generating unit's emissions are compared against the emissions limits for that unit. If the emissions are higher than the limit, then the allowance cost is applied to the excess emissions; if the emissions are lower than the limit, the allowance credit is applied to the difference. Note that emissions could be over the limit for some emission types and under for other types. Also, some units might be within their limits, while other units are over their limits. Thus, there can be both emission costs and emission credits associated with the same production costing dispatch.

At the site level, the total emissions by units assigned to a site are compared with the emissions limit for that site. If the emissions exceed the limit, then an allowance cost is applied to the excess emissions. If the emissions are lower than the limit, an allowance credit is applied to the difference. For site costs and credits, EGEAS uses the system-wide allowance cost and credit values; there are no separate site values.

At the system level, the total emissions by all the System A units are compared with the total system emission limit. If the emissions exceed the limit, then an allowance costs is applied to the excess emissions. If the emissions are lower than the limit, an allowance credit is applied to the difference. EGEAS uses the system-wide allowance cost and credit values.

In a particular analysis, you must choose only one level of application (unit, site, or system) for the emission costs and credits. These costs (and credits) are included in the production cost totals. However, no allowance costs or credits are applied to interconnected systems (B to J).

Allowance costs can also be used in conjunction with emission dispatch and constraint options (see the later section) to alter the dispatch order. For this purpose, EGEAS applies the costs to units of all areas in a multi-area run. Note that these two uses of allowance costs: for inclusion in the production cost and for altering the dispatch order, are independent. You must supply input options to tell EGEAS how the costs should be used.

Unit Emission Constraints

With the dynamic programming or pre-specified pathway methods, you can specify that EGEAS impose unit emission constraints on the production costing dispatch. During the dispatch, EGEAS reduces a unit's generation if necessary to keep the emissions within the limit you have specified for that unit. The basic steps are as follows:

1. Trial load the unit, according to its position in the economic dispatch order, and determine the unit's generation. Also calculate the emissions based on that generation.
2. If the emissions are within the limit, leave the unit loaded and proceed to the next unit.
3. If any emissions are over the limit, determine the maximum possible generation, considering all unit emission limits, and "delay" the unit to be loaded where it will generate at the desired level. The method is essentially the same as that used for limited energy plants. Refer to the Limited Energy Plants section of Chapter 3 for details on the methodology.

If the emissions limit is reached while dispatching part of a multiple-block unit, no further blocks of that unit will be dispatched. When several emission types are limited, EGEAS reduces the generation until all emissions are satisfactory. If subweekly time periods are modeled, all the allowed emissions are made available in the first subweek; the generation in later subweeks is dependent on the remaining allowed emissions after the earlier subweeks. With multi-area modeling, unit emission constraints can be applied to units of all areas.

Site Emission Constraints

Site emission constraints are applied in much the same manner as unit emission constraints. Generation by individual units is curtailed where necessary to keep the site emissions below the limit. As each unit (or block of a unit) is dispatched, its emissions are compared with the remaining amount of emissions allowed for its site. If the emissions are too high, the unit's loading is delayed to where the emissions will be within limit. EGEAS does not attempt to allocate the allowed emissions to individual units at the site; instead all of the allowed emissions are available for the first unit dispatched; later units get whatever remains.

Note that you can specify both unit and site emission constraints in the same analysis. Also, you can provide limits for multiple emission types both for individual units and for sites. When EGEAS loads a unit, it checks all the relevant emission limits and reduces the generation as much as needed to ensure that all the constraints are satisfied.

With multi-area modeling, you can define sites for each area; however, you should not include units from two different areas in the same site. To limit the total amount of emissions for one of the interconnected systems, you could include all of that system's units in a single site and impose a site emissions constraint.

System Emission Dispatch and Constraint Options

EGEAS provides three methods for modifying the dispatch based on system-wide emission totals. The single-step and iterative methods work with only one emission type (that you select), but the allowance cost method works with all emission types.

For the *single-step method*, EGEAS first performs an economic dispatch without regard to the system emissions limit. If requested, unit and/or site emission constraints are applied, however. After the dispatch, EGEAS calculates the total emissions of the selected type. If the emission limit is satisfied, production costing is complete. If the emissions are too high, then EGEAS modifies the dispatch order and performs production costing again.

For modifying the dispatch order, EGEAS calculates an emission adjuster for each unit. To get this adjuster, EGEAS determines the rate of emission production, in emission mass units per MWh, using the unit's heat rate, the emission production rate, and the emission content of the fuel. This value is then multiplied by the allowance cost value to get the emission adjuster in $/MWh. You can specify whether EGEAS should use the unit-specific or system-wide values for the allowance costs. EGEAS then re-orders the generating units, based on the sum of their fuel cost, variable O&M cost, dispatch modifier, and emission dispatch adjuster. With this revised dispatch order, EGEAS then performs production costing again.

If EGEAS is unable to satisfy the system emission limitation with the above procedure, it takes one of the following actions, depending on the option you select:

1. Allow the emissions to exceed the limit. If you have specified the application of allowance costs, the system is penalized by the inclusion of these added costs.
2. Curtail the output of generating units so that the emissions are reduced to the limit. The loss of generation means that the unserved energy is increased.
3. For a DP optimization only, mark the current state as infeasible. Thus, the DP selects an optimal expansion plan in which the system emission constraints is satisfied in every year (assuming such a plan exists).

The *iterative method* starts off in the same manner as the single-step method. EGEAS first performs an economic dispatch without regard to the system emissions limit. If the emissions limit is satisfied, the process is complete.

Otherwise, EGEAS modifies the dispatch order and performs production costing again. The iterative method differs from the single-step method in that the emission adjusters are not based on the user-supplied allowance cost values, but on internally-generated values. Essentially, EGEAS tries to produce the most economical dispatch where the emissions limit is satisfied. A high value causes the most highly-polluting units to be loaded later and hence generate less energy. But because it alters the economic dispatch order, the production cost increases. At each iteration, EGEAS tries a new allowance cost value, determines a revised dispatch order, and performs production costing. If the emissions are too high, EGEAS tries a higher allowance cost value for the next iteration to reduce the emissions. If the emissions are below the limit, EGEAS tries a lower allowance cost value for the next iteration to lower the cost. The process continues until it converges on an optimum solution. You can specify the convergence criterion (in percent) and you can also specify the maximum number of iterations.

Following is an example:

Iteration	Emission Adjuster $/Ton	Emissions Ton	Emissions Limit Ton	Over/Under Limit Ton
1	0.0	180896.2	98100.0	82796.2
2	47199.8	12007.1	98100.0	-86092.9
3	9440.0	14952.0	98100.0	-83148.0
4	1888.0	62036.7	98100.0	-36063.3
5	377.6	138863.1	98100.0	40763.1
6	1132.8	95850.6	98100.0	-2249.4
7	755.2	128710.5	98100.0	30610.5
8	944.0	102829.8	98100.0	4729.8
9	1038.4	95850.6	98100.0	-2249.4
10	991.2	102829.8	98100.0	4729.8
11	1038.4	95850.6	98100.0	-2249.4

This table shows the emission adjuster used at each iteration and the resulting emissions. For iteration 1, the emissions are far above the limit, so EGEAS chooses a high adjuster for iteration 2. The resulting emissions for iteration 2 are far below the limit, so EGEAS proceeds on subsequent iterations to select adjusters between the extreme values for the first two iterations. In the final several iterations the program is oscillating between two different values for the emissions. This is because the small differences in emission adjusters are not sufficient to change the dispatch order. Thus, for this case, as in most cases, it is not possible to get the emissions to exactly match the emissions limit. If the emissions cannot be brought within the specified limit, even with a high allowance cost value, EGEAS follows one of the three courses of action that were outlined with the single-step method.

For the *allowance cost method*, EGEAS calculates emission adjusters for each unit based on the emission allowance values that you supply. If there are multiple emission types, adjusters are calculated for each type and added together. The adjusters are then added to the marginal operating costs for each unit and the resulting sum is used for determining the dispatch order. Units with high emissions are penalized and hence will be loaded later in the dispatch order, units with low emissions are favored, relative to more polluting units, and are loaded earlier. With the allowance cost method, EGEAS does not impose a system emission constraint. If there are any emissions above the system emission limit, EGEAS simply accepts them.

Banking of Emission Allowances

Emissions banking is an optional feature that can be used to even out differences in emissions production from year to year. Low emissions in one year can help to offset high emissions in the next year. The methodology is as follows. If the total system emissions for a particular year are below the limit that you specified for that year, then the emission allowances are carried forward to the next year. These allowances are used to increase the emissions limit for that year. If the total emissions exceed the limit, however, the excess is not carried forward; but a cost can be applied to that excess. If desired, you can supply a starting balance for emission allowances, a

value that will be used to increase the limit for the first study year. At the end of the study period, an allowance credit, if you specified one, is applied to any unused allowances. Emissions banking is not performed in the extension period. Following is an example.

Year	Emissions Ton	Emissions Limit Ton	Limit with Banking Ton	Over/Under Limit Ton	Allowances Carried Forward Ton
Starting balance					**20,000**
1	110,000	100,000	120,000	-10,000	10,000
2	105,000	100,000	110,000	-5,000	5,000
3	105,000	100,000	105,000	0	0
4	110,000	100,000	100,000	+10,000	0
5	105,000	100,000	100,000	+5,000	0
6	95,000	100,000	100,000	-5,000	5,000
7	95,000	100,000	105,000	-10,000	0

Note that no allowances are carried forward at the end of the last study year (year 7).

Fuel Switching

When there are generating units with two fuels, emissions can sometimes be reduced by switching fuels, to use more of the cleaner fuel. Fuel switching can be used as the only option for reducing emissions or in conjunction with other methods such as applying emission adjusters. When used together, the fuel switching procedure is tried before the methods that alter the dispatch order with emission adjusters. The methodology is as follows:

Perform production costing with each dual-fueled unit set to the user-supplied target percentage for the secondary fuel. Apply unit and site emission constraints if they have been specified. Also, if the allowance cost method is specified for system emissions, the emission adjusters are used in determining the initial dispatch order. If the system emissions are within limit, the procedure is complete.

Switch fuels in attempt to reduce emissions. Depending on the option you select, EGEAS starts either with the first unit dispatched or the last unit and proceeds through the dispatch order in forward or reverse order. Whenever a dual-fueled unit is found, EGEAS checks whether the unit's secondary fuel percentage can be changed and if such a change will result in lower emissions. If so, the unit's secondary fuel percentage is changed, either upward or downward, as far is necessary to get the emissions within limit (or are far as possible). If the emissions are still too high, the procedure continues with the next unit. If the emissions are brought within limit, the procedure is complete.

Leave all units with the secondary fuel percentages from the previous step. If either the single-step or iterative method has been selected for system emissions, alter the dispatch order and dispatch the system again.

Note that the methodology for switching fuels is basically the same as is used with fuel switching to meet maximum and minimum fuel constraints. This methodology is described in the Dual Fuels and Fuel Constraints section of Chapter 3. In the case of conflicts between switching for

fuel constraints and switching for emission constraints, the priority is as follows: (1) satisfy the maximum fuel constraints, (2) satisfy the emission constraints, (3) satisfy the minimum fuel constraint.

Emissions Effects

EGEAS provides information on the cost impact of imposing emission constraints. Two types of emission effects information are provided. First, if you have selected the single-step or iterative method for system emission constraints, EGEAS will report the system emissions total without the emission adjusters as well as the final emissions total with the adjusters. If the emission adjusters are not needed, the two amounts are the same.

Second, EGEAS allows you to see the full impact of any type of emissions modeling. For this feature, EGEAS actually performs two complete analyses: two dynamic program optimizations or two pathway simulations. For the first analysis, no emissions constraints are imposed and the dispatch order is not altered to reduce emissions. Emissions are still calculated and emission allowance costs and credits are included in the production cost, if you request them. For the second analysis, all the emissions modeling options that you have selected are implemented. One other difference between the two analyses is that you can designate one or more planning alternatives as being emissions alternatives – alternatives whose main purpose is to reduce emissions. These emissions alternatives are included only in the second analyses. EGEAS provides reports showing the differences in emissions and in expansion plan costs between these two analyses.

Emission Reduction Constraints

Overview

Up to this point, all emission constraints have been mass-based – the constraints have been applied to an amount of emissions produced, such as tons or pounds of emissions per year. The emission reduction constraints add the ability to constrain the rate of emissions produced. Specifically, constraints can be applied to the amount of emissions produced per unit of generation. For example, a constraint can be placed on the amount of CO_2 from existing fossil-fired generating units per gigawatt hour of generation. To help lower the rate, the generation from certain designated clean units (e.g. renewables) can be included in the generation. This feature was designed to help model compliance with the Environmental Protection Agency's Clean Power Plan (CPP), but with flexibility to help with other modeling needs as well.

Emission Rate Calculation

An emission rate is defined by three key parameters:

- Emission type (e.g. CO_2)
- Target units: a group of units whose production of the designated emission type is to be reduced (e.g. fossil steam units)

- Clean units: a group of units which produce minimal or no emissions, which are included to reduce the overall emission rate (e.g. renewables such as solar or wind)

The basic formula for calculating the emission rate is as follows:

$$\text{Rate} = \frac{\sum \text{EMIS}_{\text{TARGET}} * \text{PN}_{\text{TARGET}}}{\sum \text{GEN}_{\text{TARGET}} * \text{PD}_{\text{TARGET}} + \sum \text{P}_{\text{CLEAN}} * \text{GEN}_{\text{CLEAN}}}$$

Where:

Rate	=	emission rate
$\text{EMIS}_{\text{TARGET}}$	=	emissions for a target unit
$\text{GEN}_{\text{TARGET}}$	=	generation for a target unit
$\text{GEN}_{\text{CLEAN}}$	=	generation for a clean unit
P_{CLEAN}	=	fraction of a clean unit's generation to use (0 to +1)
$\text{PN}_{\text{TARGET}}$	=	fraction of a target unit's emissions to use (-1 to +1)
$\text{PD}_{\text{TARGET}}$	=	fraction of a target unit's generation to use (-1 to +1)

The emission rate is in units of emission per gigawatt hour. For example, tons CO_2 per Gwh. If a rate in lb/Mwh is desired, simply multiply the ton/Gwh value by 2.

The target units are included in both the numerator (emissions) and denominator (generation). The clean units are included only in the denominator. Even if a clean unit produces emissions of the specified type, such emissions are ignored. The user must select the units to include in each constraint. There must be at least one target unit and there is no limit on the number that can be included. Clean units are optional, but are clearly essential to reducing the rate. If there are no clean units and the emission production rates of all the target units are above the constraint limit, there is no way the constraint can be satisfied, except by not operating any of the target units.

Three parameters are included for specifying the portion of units' generation or emissions to include in the rate calculation: $\text{PN}_{\text{TARGET}}$ and $\text{PD}_{\text{TARGET}}$ specify the portion of the target unit emissions and generation, respectively, to include in the calculation. These fractions provide flexibility for modeling various regulatory requirements or special circumstances. The P_{CLEAN} parameter specifies the portion of the clean unit's generation to include in the calculation. This parameter can be used if a given unit is being used to contribute towards helping meet more than one constraint, or for other special circumstances. Note that these parameters are specified individually for each target and clean unit; i.e. one unit might contribute 100% and another only 50%.

For modeling flexibility, the emission and generation contribution of a target unit can be either positive or negative. A negative emission contribution means that increased generation by such a unit would actually lower the rate. The only restriction is that the fraction of a unit's emissions to use is between -1 and +1. In the unlikely event that too many units were assigned negative emission contributions, the emission rate could conceivably become negative. In such a case, the rate would be considered to be zero and the constraint would automatically be satisfied.

A negative generation contribution means that increased generation by that unit could greatly increase the rate. In the unlikely event that there is too much negative generation contribution, the denominator could become zero or negative. In such a case, the rate is treated as being undefined and the constraint is automatically not satisfied.

Constraint Methodology

EGEAS uses two different methods for applying emission reduction constraints. The Adjust Dispatch Method alters the production costing dispatch order so that generation by target units is reduced. The Reject State Method works with a dynamic program optimization to reject states where the emission rate is above the constraint limit. Either method can be used by itself or they can be used together. It is also possible to use neither of the methods, but simply allow EGEAS to calculate the emission rate for information purposes only.

Adjust Dispatch Method

The first method for applying emission reduction constraints involves reducing the generation by target units by altering the production cost dispatch. The dispatch is altered only to the extent needed to get the emission rate below the constraint limit. An iterative process is used.

The method starts with a production cost dispatch with no provision made for the emission reduction constraints. After the dispatch, the program checks all of the constraints. If they have all been met, no additional steps are needed. The process also stops if there is any constraint where all the target unit have emission production rates that exceed the constraint limit and there are no clean units. In this case, there is nothing that can be done to bring the overall rate below the constraint limit.

Otherwise, the program attempts to alter the dispatch to reduce the emission rate. The goal is to reduce the generation by target units and/or increase the generation by clean units, so the resulting emission rate is within the limit. An iterative process is used to determine the extent of alterations needed to achieve the desired result.

To alter the production cost dispatch order, the program adds an amount to the marginal cost for each of the target units. This effectively raises the cost used to determine the dispatch order so the target units are dispatched later and hence, generate less energy. The marginal cost adder for each target unit is based on an emission cost, in units of dollars per unit of emission, so that units that generate more emissions are penalized more heavily. If less than the unit's full emissions are used in the rate calculation (see Emission Rate Calculation section above), then the marginal cost adder is reduced accordingly. Also, if the contribution is negative, the adder is set to zero because increasing generation by such a unit actually lowers the rate. As an example, the marginal cost adder for a unit with positive contribution to CO_2 emission is calculated as:

Marginal Cost Adder ($ / Mwh) =
emission adjuster ($ / ton CO_2) *
emission production rate for unit (ton CO_2 / Mwh) *
fraction of unit's emission used in rate calculation

As mentioned above, the first production cost run (iteration 1) has no emission adjuster (effectively an adjuster value of zero). The second iteration uses a very high value for the adjuster, such that the dispatch order is almost entirely based on the amount of emissions produced. This effectively drives the emissions (and emission rate) as low as possible. If the emissions rate is still above the constraint limit, then the constraint cannot be satisfied by changing the dispatch order and the iterative process stops.
If the constraint was satisfied at the second iteration, the objective is partially achieved. However, the large additions to the marginal costs may have driven the production cost unnecessarily high. What remains is to find a smaller value of the emission adjuster which still results in the emission rate within the constraint limit. Subsequent iterations use different values for the emission adjuster until the program finds the lowest adjuster value where the constraints are met. The process continues until convergence is achieved or the maximum number of iterations is reached. Convergence is assumed if the emission adjusters for the last two iterations are within a specified percent of the larger adjuster.

Reject State Method

With a dynamic program run, the user can specify one or more planning alternatives that represent clean units, such as renewables. These units can be included in one or more rate-based constraints. Existing and committed units can also be designated as clean units. Note that if multiple units of a given planning alternative are installed, all such units will be included in the same rate-based emission constraints, if any. If a given DP state includes such units, then the emission rate will be reduced. If the emission rate for a state exceeds the limit for that year, then that state will be marked as infeasible and excluded from the optimization. Thus, the DP will find the least cost plan, if any, that satisfies all the rate-based emission constraints.
If the adjust dispatch method is also selected, then the feasibility of a state will be based on the results after the production costing has been adjusted to try to meet the rate constraints; i.e. after the iterative process has completed. If the production cost method is not selected, then a given DP state is feasible only if all the emission reduction constraints are satisfied without any dispatch adjustments.

Annual vs. Segment-Level Constraints

Emission reduction constraints can be applied at either an annual or segment level. For an annual constraint, the emission rate is calculated using the total annual generation and emissions for each of the generating units included in the constraint. If the rate is less than the annual constraint limit, the constraint is satisfied.

For a segment level constraint, the program applies a constraint in each segment of the year. Thus, if 12 segments have been defined, approximating monthly periods, then there are 12 constraints. For each segment, the emission rate is calculated using the generation and emission values for that segment. The segment constraint limits can be different for each segment if desired. The overall emission reduction constraint is considered to be satisfied if the constraint is satisfied in every segment. Thus, even if the constraint limit

is the same in every segment, the segment level represents a tighter constraint than the annual level because there are more conditions that must be met.

The application of segment constraints with the reject state method is straightforward. For each state, the emission rate is calculated for each segment. If all the segment constraints are satisfied, the state is feasible.

The application with the adjust dispatch method is more involved. At each iteration, the program performs productions costing for each segment and determines the emission rate for the segment. Different emission adjusters and marginal emission adders may be used in each segment. If the rate is below the limit in a particular segment, then a smaller emission adjuster is used for the next iteration – to try to find a lower cost solution. If the rate is above the limit, then a larger adjuster is used for the next iteration – to try to get the emission rate below the limit. The iterative process continues until convergence is achieved in all segments or the maximum number of iterations is achieved. If convergence is achieved in some segments before it is achieved in others, the process still continues in all the segments.

Multiple Constraints

It is possible to apply multiple emission reduction constraints in the same EGEAS run. Individual constraints can have different emission types and some can be applied at an annual level while others are at a segment level. Some can use the adjust dispatch method while others use the reject state method, and still others use both methods. However, there are a number of issues to be addressed with multiple constraints.

Applying multiple constraints is pretty straightforward with the reject state method. After production costing has been performed for all the subperiods and potentially with multiple iterations for the adjust dispatch method, the program determines whether each constraint has been satisfied. If even one constraint is not satisfied, the state is rejected as infeasible.

The alter dispatch method involves an iterative process to modify the dispatch order to bring the emission rate within the constraint limit. When the adjust dispatch option is selected for multiple constrains, all the constraints use the same set of iterations. There are two reasons for this. First, having multiple sets of iterations, particularly if one set is nested inside another, results in greatly increased running time. Second, and more importantly, changes made in the dispatch order for one constraint can impact the results for other constraints

For the first iteration, there are no marginal cost adders for any of the constraints and the production cost is performed for all the subperiods. If all constraints are satisfied, both annual and segment constraints, the process is done. Otherwise, a large emission adjuster is selected for each constraint. Marginal cost adders are determined for all the units. When a unit has adders from more than one constraint, the program uses the largest one. It does not matter if the adders are from constraints with the same or different emission types. It never uses the sum of the adders. For the second iteration, the production cost is performed for all the subperiods again. If any constraint is not satisfied, the process is terminated – the constraints cannot all be satisfied.

Otherwise, the program proceeds through additional iterations. At each iteration, a new emission adjuster is determined for each constraint. If the emission rate was too high for a constraint at the previous iteration, the adjuster for that constraint is increased. If the rate was below the constraint limit, the adjuster for that constraint is decreased. For a segment-level constraint, the new emission adjusters may be higher in some segments and lower in others. The process continues until all the constraints have converged to a solution or the maximum number of iterations is reached.

Practical Considerations and Warnings

When multiple rate-based emission constraints are applied, they can interfere with each other. Changes in the dispatch order for one constraint may affect the results with other constraints. The results may oscillate from iteration to iteration and the process may take significantly more iterations to converge or may not converge completely to an optimal solution. When multiple constraints are applied, the user should check the results carefully, including the log report in the CANAL output, to verify that the final results are reasonable.

For each constraint, the iterative process uses lower and upper bounds on the emission adjuster. As the iterations progress, these bounds are updated until the lower and upper bounds converge. However, with multiple constraints, later iterations sometimes prove that the bounds from earlier iterations were too tight and the true optimal solution is outside those bounds. To help deal with this situation, the program will widen the bounds when necessary to increase the chance of finding the best solution. The method is not fool-proof, but this added step does tend to improve the results in many cases.

A bisection algorithm is used to select the dispatch adjuster for the next iteration – essentially setting the new adjuster to the average of the current lower and upper limits. However, to shift the results slightly toward the feasible end, the new adjuster is set at 45% of the lower limit + 55% of the upper limit. (Higher adjusters result in lower emission rates).

When must-run units are designated as target units, the adjust dispatch option has no impact on their generation – the units continue to operate to their fullest capability because of their must-run status. Using an emission adjuster to penalize the units and cause them to generate less energy has no effect. Depending on the overall mix of target units, increasing the emission adjuster can even increase the emission rate instead of lower it. Thus, if a significant amount of target unit generation comes from must-run units, the adjust dispatch option is not effective. In this case, the focus should be on the reject state option.

Because there is just one iterative loop that serves all the constraints, some constraints may go through an iterative process even though they were satisfied at the first iteration. This is needed because changes for one constraint may change another constraint from being feasible to being infeasible. If the impact of other constraints is minimal, the results for the constraint may end up essentially where they started. There is one case where iterations are not needed for a constraint: if all the constraints are applied at the segment level (i.e. no annual level constraints) and all the constraints in a particular segment are satisfied at iteration 1, then no emission adjusters are needed. The program still has to perform the dispatch for each iteration

Limitations

The following limitations apply to emission reduction constraints:

- If there are multiple emission reduction constraints, the program does not automatically assign clean units to specific constraints based on need. The user must designate in advance which clean units are to be used for which constraints. A clean unit can be shared by multiple constraints, but the user must specify the portion assigned to each constraint.
- For segment-level constraint, the user must specify the constraint limit for each segment. The limits can be the same in each segment or they can be different. The program does not provide an automatic way of determining the segment limits.
- With multi-area modeling, constraints can be applied only with System A. Units from other systems cannot be included in any constraints.
- Emission rate calculations and constraints are allowed only with dynamic programming and pre-specified pathway option.

Emission reduction constraints are not allowed when any of the following are being modeled:

- System emission constraint with either the 2-step option or the iterative option
- Switching fuel on dual-fueled units to lower emissions
- Benders decomposition optimization

Emission reduction constraints are not allowed for any emission type where any of the following features are being modeled:

- Emission allowance costs to alter the production cost dispatch
- Emission allowance costs or credits added the production cost
- Unit emission constraint on any generating unit that is a target or clean unit for the emission reduction constraint
- Site emission constraint where any generating unit that is at that site is also a target or clean unit for the emission reduction constraint

Cautions

The following modeling features might interfere with meeting rate-based emission constraints or might themselves be impacted by the rate-based constraints

- Minimum fuel constraints. Placing a lower limit on fuel used by target units could make it difficult to lower the generation by target units and thus meet the rate-based emission constraints.

- Maximum fuel constraints. Placing an upper limit on fuel used by clean units could make it difficult to get the necessary generation to reduce emission rates.
- Fuel switching for any fuels used by target or clean units could complicate the iterative process and possibly result in non-convergence.
- Storage optimization. Reducing generation by fossil steam units could end up reducing energy available for charging storage units or reduce the benefits of storage operation due to using more costly generation for charging.
- Multi-area modeling with tie-line limits. The economics associated with multi-area modeling are based on actual costs, and so would not be impacted by the cost adders used to alter the dispatch order. However, the cost adders could impact the amount of energy that is potentially sold or bought.
- If a significant proportion of the target units for a constraint are must-run units, then the adjust dispatch method may not converge. In fact, using a larger emission adjuster may lead to a larger rather than a smaller emission rate. The reason is that the adjust dispatch method aims to reduce generation by target units by placing them later in the dispatch order. However, must-run units remain at the beginning of the dispatch order because of the must-run requirement. In this case, it is better to turn off the adjust dispatch option and just use the reject state method for this constraint.

10
ADDITIONAL CAPABILITIES

Overview

In addition to the basic modeling capabilities described in the previous chapters, EGEAS can handle several additional features such as uncertainty analyses, avoided cost calculations, and price elasticity modeling. This chapter describes these specific features of EGEAS in more detail.

Uncertainties Modeling

Automatic Sensitivity Analysis

The automatic sensitivity analysis feature in EGEAS allows systematic evaluation of how changes in the values of certain input variables affect the optimal or pre-specified expansion plan. Different values for input parameters may reflect either lack of information (e.g., efficiency or availability of a resource) or uncertainty about future realizations of forecasted input variables (e.g., load growth rates, cost escalation, construction lead times, licensing). The effects on the expansion plan are analyzed by monitoring not only changes in optimal capacity additions but also the effects on costs, fuel use, emissions, reliability, incremental costs for meeting reliability targets, and other attributes of interest.

The EGEAS sensitivity analysis feature allows you to specify variations in one or more input parameters. Variations in multiple parameters can be independent (all possible combinations) or joint (parameters change in tandem). For each possible combination of the input parameter values, EGEAS performs the requested analysis (e.g., dynamic programming, Benders' decomposition, pre-specified pathway). Thus, the sensitivity analysis is automated, saving you time and reducing the chance for errors.

There are three options for specifying which combinations of input parameter values should be evaluated. To illustrate these options, consider three input variables: x, y, and z. Assume that each variable assumes 2 values: x_1, x_2, y_1, y_2, z_1, and z_2. For joint variation, the user simply lists the desired combinations as, for example:

Scenario 1 x_1, y_1, z_2

Scenario 2 x_2, y_1, z_2

Scenario 3 x_1, y_2, z_1

For independent variation, EGEAS tries all possible combinations, which for the current example would produce:

Scenario 1 x_1, y_1, z_1

Scenario 2 x_1, y_1, z_2

Scenario 3 x_1, y_2, z_1

Scenario 4 x_1, y_2, z_2

Scenario 5 x_2, y_1, z_1

Scenario 6 x2, y_1, z_2

Scenario 7 x_2, y_2, z_1

Scenario 8 x_2, y_2, z_2

For partially-joint variation, you can link the variations in two or more variables and let EGEAS try all possible combinations with the remaining variables. If variables x and y in the above example are linked, the possible scenarios are:

Scenario 1 (x_1, y_1), z_1

Scenario 2 (x_1, y_1), z_2

Scenario 3 (x_2, y_2), z_1

Scenario 4 (x_2, y_2), z_2

Accounting for joint variation of input parameters allows for modeling the interdependence of certain quantities in real life. For example, if fuel costs and load growth are investigated, scenarios with high fuel costs and high load growth would be considered improbable. Such a combination of parameters could be excluded from the list of possible scenarios.

Avoided Costs

Avoided costs measure the expected change in total costs due to an increase or decrease in load or demand. For use in an IRP, avoided costs should be calculated from a long-run perspective. Long-run avoided costs include both the variable costs associated with additional output and the fixed capital investment required to maintain reliable service. EGEAS provides an automatic feature for calculating the avoided costs. The avoided cost values are stored by EGEAS and can be imported into DSM screening models such as SRC COMPASS or EPRI's DSManager. The avoided costs are then used to determine the benefits of implementing a DSM measure.

In order to obtain the avoided cost information, EGEAS requires that you provide several input parameters. First, you specify (input to the EDIT program) one or more resources as avoided cost resources with a designated flag in the basic resource data. Only existing and committed resources can be designated as avoided cost resources. Once the EGEAS data base is created with information as to which resources are designated as avoided cost resources, you can instruct EGEAS to calculate avoided costs in the CANAL run. This feature is specified as a marginal analysis option and it is allowed with Bender's decomposition, the dynamic program and the pre-specified pathway. After the CANAL run is complete, you can then request the REPORT program to produce a detailed report on the avoided cost information.

When the avoided cost calculation option is turned on, EGEAS will perform the expansion plan optimization or simulation twice. EGEAS first performs the analysis with the inclusion of the avoided cost resources; this is the alternative case. Then the analysis is performed without the avoided cost resources; this is the base case. Note that when EGEAS outputs the final expansion plan(s), complete information is saved only for the base case (without the avoided cost resources); only summary information (e.g., total capital and operating costs) is saved for the alternate case (with the avoided cost resources). Therefore, EGEAS can produce reports only for

plans derived from the base case. However, EGEAS can produce a detailed Avoided Cost Report, which shows the major differences between the two cases. Avoided costs are computed for each subperiod (if subperiods are modeled) as well as for each year; cost values are reported in both current year dollars and present worth dollars (discounted to the base year).

There are two components of avoided costs that can be calculated in EGEAS: the capacity (demand) costs and the energy costs.

Avoided Capacity Costs

Avoided capacity or demand costs are the costs of building and maintaining new generating plants and purchasing power to meet increases in electric demand. Hence the cost values are expressed in units of dollars per kilowatt of capacity for the added capacities of the avoided cost resources.

EGEAS first calculates the avoided capacity costs on an annual basis. As mentioned previously, EGEAS evaluates the alternative and base cases (with and without the avoided cost resources, respectively). In each case, the calculation of the capacity costs includes capital fixed charges and fixed O&M costs. If detailed costs are modeled, the calculation includes the sum of the detailed capital and fixed O&M costs over all cost components. The difference ($) between the two cases (i.e., base case - alternative case) is then divided by the added rated capacity (MW) of the avoided cost resources to obtain the annual avoided capacity cost ($/MW).

If subperiods are modeled, EGEAS provides the following three options for allocating the annual capacity cost to individual subperiods:

1. With Option 1, EGEAS allocates the annual capacity cost to individual subperiods in proportion to the number of hours in each period.
2. With Option 2, the annual capacity cost is allocated to subperiods in proportion to the loss of load hours in each period. Note that if LOLP is calculated only on an annual basis (input option in CANAL), the LOLP in each period is assumed to be the same; this option will yield the same results as Option 1.
3. With option 3, the entire annual capacity cost is allocated to the period with the highest LOLP. The capacity cost in other subperiods is set in proportion to the loss of load probability. With this option, the subperiod capacity costs do not sum to the annual capacity cost. If LOLP is calculated only on an annual basis (input option in CANAL), the LOLP in each period is assumed to be the same; in this case, the entire capacity cost is allocated to every subperiod.

The following example illustrates how EGEAS allocates the capacity costs to individual subperiods.

Assume:

Segment	Subweek	Hours	LOLP	LOLH
1	1	780	.0030	2.340
	2	1,404	.0020	2.808
2	1	1,560	.0025	3.900
	2	2,808	.0015	4.212
3	1	780	.0020	1.560
	2	1,404	.0015	2.106
Annual		8,736	.0019	16.926
Annual capacity avoided cost = 1,000. ($/kW)				

With Option 1, the subperiod capacity costs are:

Segment	Subweek	Capacity Cost ($/kW)
1	1	(780/8736) x 1,000 = 89.3
	2	(1404/8736) x 1,000 = 160.7
2	1	(1560/8736) x 1,000 = 178.6
	2	(2808/8736) x 1,000 = 321.4
3	1	(780/8736) x 1,000 = 89.3
	2	(1404/8736) x 1,000 = 160.7

With Option 2, the subperiod capacity costs are:

Segment	Subweek	Capacity Cost ($/kW)
1	1	(2.340/16.926) x 1,000 = 138.4
	2	(2.808/16.926) x 1,000 = 165.9
2	1	(3.900/16.926) x 1,000 = 230.4
	2	(4.212/16.926) x 1,000 = 248.8
3	1	(1.560/16.926) x 1,000 = 92.2
	2	(2.106/16.926) x 1,000 = 124.4

With Option 3, the subperiod capacity costs are:

Segment	Subweek	Capacity Cost ($/kW)
1	1	(0.0030/0.0030) x 1,000 = 1,000.0
	2	(0.0020/0.0030) x 1,000 = 666.7
2	1	(0.0025/0.0030) x 1,000 = 833.3
	2	(0.0015/0.0030) x 1,000 = 500.0
3	1	(0.0020/0.0030) x 1,000 = 666.7
	2	(0.0015/0.0030) x 1,000 = 500.0

Avoided capacity costs depend on the utility type (the member-based utility, G&Ts, power supplies, etc.). If a utility does not require capacity in the next 20 years, the avoided capacity cost may be zero. However, if additional capacity is required in the next 20 years, then the simplest way to estimate avoided capacity costs is to develop a proxy for these costs. The proxy cost is determined from a supply-side resource with the lowest investment cost, usually a combustion turbine. The logic behind this approach is that avoided capacity costs are synonymous with avoided reliability costs. If a system planner were to require a capacity addition for the sole purpose of maintaining reliability, fuel savings would not be considered. The chosen plant would be the one with the lowest construction cost, which in most cases turns out to be a combustion turbine.

Avoided Energy Costs

By definition, energy costs represent the costs that vary with the level of energy produced. Avoided energy costs represent the additional fuel and variable operating costs or energy cost component of a contract required to produce certain amount of energy. Hence the energy cost values are expressed in units of dollars per megawatt hour of generation for the added energy from the avoided cost resources.

Avoided energy costs can be estimated by evaluating the day-to-day operation of a utility system, incorporating fuel costs, variable O&M expenses, unit efficiency data, maintenance schedules, and availability factors. The costs also depend on the types of purchased power contracts and generating resources available to the utility.

To obtain the avoided energy costs, EGEAS evaluates the alternative and base cases (with and without the avoided cost resources, respectively). In each case, EGEAS sums the components of the energy costs (fuel costs + variable O&M costs + purchased energy costs + unserved energy costs + emission costs - emission credits). If detailed costs are modeled, the calculation includes the sum of the detailed variable O&M costs over all cost components. The difference ($) between the two cases (i.e., base case - alternative case) is then divided by the added energy (MWh) from the avoided cost resources to obtain the avoided energy cost ($/MWh). If subperiods are modeled, EGEAS computes the subperiod energy costs based on the operating costs derived for each subperiod.

Hourly Avoided Capacity Cost File

Avoided costs are often used for evaluating the benefits of implementing a DSM measure. To accommodate such evaluation, EGEAS can output the avoided capacity cost data to a file which can then be used by other DSM models.

To obtain the Hourly Avoided Capacity Cost File, you must first let EGEAS create a data base with the designated avoided cost resources and instruct EGEAS to calculate the avoided costs. When requesting the file, you must also choose one of three options which specify how avoided capacity costs should be calculated and allocated to individual subperiods. Input requirements and the three calculation options (including sample calculations) have been described in detail in the **Avoided Capacity Costs** section. Note that the Hourly Avoided Capacity Cost File is a file produced by EGEAS which contains the hourly avoided cost values in an hourly file format; it should not be confused with the Avoided Cost Report which you may request as a printed output from EGEAS. Since the two are independent, you may choose different capacity cost calculation options for the hourly file and the printed report.

The following describes the method used in EGEAS to derive the hourly capacity cost values and output them to the file. EGEAS first determines the avoided capacity costs for individual subperiods in the same manner as previously described. When subperiods are modeled, EGEAS knows which hours of the year are included in each subperiod. Use the same example given in the **Avoided Capacity Costs** section and assume that, for a given year, the first hour on the sixth day of January belongs to segment 1, subweek 2. In this example, there are 1,404 hours within a year that belong to this subperiod. EGEAS will then divide the appropriate capacity cost for this particular subperiod (i.e., segment 1, subweek 2) by 1,404 and assign the resulting value to the first hour on January 6th and to all of the other hours that belong to same subperiod. Depending on the specified capacity cost calculation option, the hourly value for those hours in this example would be:

For Option 1: (160.7/1404) = 0.1145 $/kW

For Option 2: (165.9/1404) = 0.1182 $/kW

For Option 3: (666.7/1404) = 0.4749 $/kW

Incremental Costs

In EGEAS, incremental cost (or marginal cost) is defined as the cost of generating an additional unit of energy at a specific load level. EGEAS calculates incremental costs based on marginal operating costs (fuel plus variable O&M) for the resources, fixed O&M costs are not included. Unit forced outages are treated probabilistically and all resources are assumed to be independent; resource characteristics are derated to reflect maintenance requirements. The methodology is essentially consistent with the probabilistic production costing method used in EGEAS.

To calculate incremental costs, EGEAS convolves together the forced outages of the resources to obtain a cumulative outages distribution. The convolution process is the same as that used for production costing except that the system load is not included. Resources are convolved in the same order as they were loaded for production costing. For a given megawatt value, the cumulative outage distribution gives the probability that the capacity forced out exceeds that value. From this distribution, EGEAS determines the probability that a given resource is the marginal resource at a particular load level. The marginal resource is the last resource loaded; more importantly, it is the resource that will provide increased generation if the load is increased.

The incremental cost at a particular load level is the weighted average cost of all the resources that can be marginal at that point; the weights are the probabilities that the resources are marginal. For a thermal resource (or a hydro resource whose loading is not delayed because of

energy limitations), EGEAS uses the marginal operating cost of the resource itself. For a delayed hydro resource or the generating side of a storage resource, EGEAS uses the marginal operating cost of the thermal resource being "split"; if the load level is increased, this thermal resource will provide increased generation because the delayed unit is already at its energy limit.

As in the production costing dispatch, storage charging resources are convolved before any generating resources to reflect the need for generation to charge the storage resources. The storage charging sides are then deconvolved once the required charging energy has been supplied. Because of the charging load, the effect of storage resources is to raise the incremental cost at low load levels. However, the generation by storage resources is less expensive than that of the displaced peaking resources, so the incremental cost at the higher load levels is reduced. The incremental cost calculations are consistent with deterministic (but not probabilistic) allocation of storage energies. (The allocation options are described in Chapter 3.)

EGEAS does not consider non-dispatchable technology (NDT) resources in the incremental cost calculations. Because the output of an NDT resource is fixed (hourly values are supplied by the user), an NDT resource cannot increase its output to respond to an increase in load; thus, it cannot be a marginal generator. You can interpret the incremental cost at a load value of X megawatts in one of two ways: as the incremental cost when the system demand is X (assuming no NDT resources), or as the incremental cost when the net load (after modification for NDT resources) is X.

At any load level, the available capacity can be less than the system demand because of outages. In such a situation, an increase in load simply results in more unserved energy. Thus, the incremental cost calculation includes a term for unserved energy, with the cost of unserved energy input to the model.

Incremental costs are calculated only when a particular expansion plan is being simulated, not during expansion optimization. Costs can be calculated either for whole years or for individual subperiods for which production costing is performed. If multi-areas are modeled, incremental costs can be calculated for each independent system as well as for the total pool. EGEAS outputs incremental costs in tabular form; the cost curves are not graphed. The incremental cost tables include an equally-spaced set of points ranging from the minimum load to the peak system load or to the total system capacity. You may specify either the number of points or the megawatt increment between adjacent points. A sample incremental cost table is shown in Table 10-1. Note that EGEAS calculates the expected percent of the time that the system is at each of the load levels.

Table 10-1
Incremental Cost Table

Incremental Cost MW	Percent $/MWh	Percent of Time
1240	23.12	2.16
1390	24.64	2.75
1540	26.56	4.37
1690	28.30	5.95
1840	29.23	7.41
1990	29.86	8.25
2140	31.05	8.40
2290	35.05	8.15
2440	42.35	7.62
2590	51.70	7.17
2740	68.96	6.85
2890	70.15	6.74
3040	71.72	6.15
3190	73.72	5.56
3340	76.15	4.71
3490	79.64	3.31
3640	82.15	2.26
3790	85.03	1.38
3940	88.24	0.81
4090	93.05	0.00

EGEAS can output incremental cost tables to a special report file in tabular form suitable for printing or to a comma-delimited file suitable for input to another program or spreadsheet. Both output files are generated by the CANAL program.

Hourly Incremental Cost Files

As described in the previous section, EGEAS can calculate the incremental costs at specific load levels and provide the cost curves in tabular form for a set of points. To correlate the incremental cost to the actual load on an hour-by-hour basis, EGEAS also provides an option to output the incremental costs to a file in an hourly file format; the cost for each hour is based on the net load for that hour. EGEAS can produce an Hourly Incremental Cost File for System A, as well as for each of the interconnected systems and the pool (if multi-areas are modeled). The Hourly Incremental Cost Files are generated by the REPORT program.

To obtain the Hourly Incremental Cost File(s), EGEAS requires that you supply Hourly Net Loads File(s) for the corresponding system(s) as input to the REPORT program. The

Net Loads Files are output files which can only be generated by ORTHOG; they contain hourly loads which reflect all modifications to fit user-supplied peak loads and energy, and modifications for purchase and sale contracts (System A only). The next requirement is that EGEAS must calculate the incremental costs for the system(s) desired. You specify this feature (as a pre-specified pathway analysis option) in the CANAL run.

When hourly incremental cost files are desired, EGEAS simply uses the net loads data to determine the incremental cost for each hour based on the net load for that hour. Since EGEAS only calculates the incremental costs for a pre-specified set of points (such as those tabulated in Table 10-1), the costs corresponding to the actual load levels are determined by interpolation using the cost curves. Use the data in Table 10-1 as an example. If the net load for a particular hour were 1,315 MW (midpoint between the 1,240 and 1,390 load levels), then the incremental cost for that hour would be 23.88 $/MWh (midpoint between the 23.12 and 24.64 cost values).

Price Elasticity and Feedback Mechanisms

Long-range generation planning depends on the forecast of customer demands for electricity over a twenty- to thirty-year period. However, the demand for electricity depends to some extent on the price that the customers will have to pay for that electricity. And these prices depend on the expansion plan – what generating units or demand-side programs will be installed – and the costs of operating the system. EGEAS contains a price elasticity feature to model this feedback mechanism.

When modeling price elasticity, EGEAS makes changes in the system demand based on changes in the system average rate relative to the rate of inflation. You must specify the rate of inflation, either as a constant value for all years, or as a trajectory of rates that vary from year to year. When the system average rate increases faster than inflation, then the system demand is lowered; when the system average rate increases more slowly than inflation, then the system demand is increased. Changes in system demand for a particular year depend on cumulative changes in the system average rate up through the previous year.

To specify how the demand changes with respect to system rate changes, you must specify an elasticity parameter. With an elasticity parameter of 1.0, a 2-percent increase in the system rate above inflation would result in a 2-percent decrease in the system demand. More generally, with an elasticity parameter of r, a 2-percent increase in the system rate above inflation would results in a r * 2-percent decrease in the system demand. You can specify either a constant elasticity value or change the value from year to year. When changing the system demand, all the loads are scaled by the same amount; i.e., there is no change in the load shape, just in the annual peak load.

If you have supplied EGEAS with multiple load components or classes (e.g., residential, commercial, etc.), you can specify a separate elasticity parameter for each load class. In this case, EGEAS modifies the demand for each load class separately and then aggregates the class loads to get the overall system demand.

To illustrate how EGEAS models elasticity, consider the following example with just one load class:

Study period: 1995 to 1999		
Peak Load:	**Year**	**Peak Load**
	1995	1,000 MW
	1996	1,050
	1997	1,102
	1998	1,158
	1,999	1,216
Inflation rate: constant at 3 percent per year		
Elasticity parameter: constant at 0.5		

Assume that EGEAS calculates the following system average rates:

Year	System Average Rate	Cumulative % Increase in Rate	Cumulative % Increase in Inflation	% Increase Relative to Inflation
1995	70.00	–	–	–
1996	70.70	1.00	3.00	-1.94
1997	73.53	5.04	6.09	-0.99
1998	77.20	10.29	9.27	+0.93
1999	81.06	15.81	12.55	+2.89

In the above table, the first column gives the year and the second column gives the system average rate calculated by EGEAS. The third column shows the cumulative percent increase in system average rate from the first study year. The fourth column gives the corresponding increase in inflation. The final column shows the percentage increase in system average rate relative to inflation. For 1996 and 1997, the system average rate has increased since 1995 at a slower pace than inflation; for 1998 and 1999, the rate has increased at a faster pace than inflation. Based on this information, EGEAS would make the following changes in the system demand:

Year	Original Peak Load	% Change in Peak Load	Modified Peak Load
1995	1,000	–	1,000
1996	1,050	–	1,050
1997	1,102	+0.97	1,113
1998	1,158	+0.49	1,163
1999	1,216	-0.47	1,210

Note that elasticity never affects the load in the first two study period years. The change in the third year is based on the increase in rate from year one to year two. In this example, the percentage increase in system average rate relative to inflation for 1996 was - 1.94 percent.

Applying the elasticity of 0.5 to this value and changing the sign yields a positive increase of 0.97 percent in the peak load. For 1998, the relative increase in the system rate is +0.93 percent. Applying the elasticity parameter yields a decrease in the peak load of 0.47 percent. This decrease is applied to the original peak load of 1,216 MW to obtain the new peak load of 1,210 MW.

In its processing, EGEAS cannot simply proceed through the study period and calculate the system average rates, and then apply the elasticity parameter to change the loads. This is because the changes in loads for a year affects the system average rate calculation for that year. Thus, at the start of each year, EGEAS determines the modified peak load based on changes in system average rate up through the preceding year, performs production costing and determines costs for the current year, and then calculates the system average rate for the current year.

This order of processing poses some potential problems. First, the method breaks down if the system average rate for one year is impacted by load changes in later years. This can happen if any portion of the construction work in progress (CWIP) is allowed in rate base. In this case, a unit installed in one year contributes to the system average rate in the preceding years. This is not a problem when simply simulating an expansion plan with the pre-specified pathway, but it can interfere with the optimization performed by the dynamic program. Thus, the following limitations apply to modeling price elasticity with the DP:

CWIP should not be allowed in rate base.

Backward dynamic programming should not be performed (EGEAS needs to know the system average rate for earlier years to determine the load for later years).

Levelized system average rate should not be selected as the objective function (requires knowing the system demand for future years).

Elasticity modeling is not allowed with Benders' decomposition.

Bid-Based Pricing

EGEAS can be used for marginal pricing within bid-based pools with the unit dispatch based on bid prices. The revenue earned by each generating unit is calculated and compared against its cost to develop the net margin for the unit. A negative net margin would indicate that the unit is a likely retirement candidate.

Unit revenues are based on the market clearing price concept. These revenues are calculated during production costing, based on bid prices and the number of hours each generating unit is the marginal unit. The bid price for each unit is the marginal operating cost multiplied by the user-specified bid multiplier. If no bid multiplier is supplied, EGEAS uses a default value of 1.0. The marginal operating cost, in dollars per megawatt hours, includes the following components:

Fuel

Variable operating and maintenance cost

Detailed costs modeled as variable O&M

Emission allowance costs, if emission adjusters are used to modify the dispatch order

Dispatch modifier, if modeled

EGEAS sets the dispatch order based on the bid prices. For multiple-block units, EGEAS uses a levelized cost for each of the blocks. Must-run blocks are treated as having a marginal cost of zero; higher blocks of must-run units use the actual, rather than levelized, cost. For hours where there is insufficient capacity (loss of load hours), EGEAS uses the cost of unserved energy as the marginal price.

The procedure for determining revenues with thermal units is as follows. After each thermal unit is loaded, EGEAS determines the amount of time that the unit is the marginal unit. This time is determined by taking the difference in loss of load hours (LOLH) before and after the unit is loaded. Revenues for the energy generated when the unit is on the margin are based on the unit's bid price. The remaining revenues depend on the bid prices of units loaded later.

As each thermal unit is loaded, EGEAS determines contributions to the revenues for thermal units that were loaded previously. The price is the bid price for the current unit. The amount of energy to which it is applied for a prior unit is the product of the prior unit's operating capacity, its availability, and the hours that the current unit is on margin.

After all thermal units have been dispatched, there may still be some energy from each unit for which no revenues have been calculated. These energies are for the time when the total capacity falls short of the system demand. For these energies, EGEAS uses the unserved energy price.

Here is a simple example:

Hours in the period	168
Peak load	500 MW
Unit capacities	100 MW
Unit availabilities	1.0 (for simplicity)

Unit	$/MWH	Generation		Hours on Margin
1	10	16,800		0
2	20	16,000		18
3	30	11,800		70
4	40	6,600		70
Unserved	50	400		10
		51,600		168

The revenues would be calculated as follows:

Unit 1:		(20 $/MWH) * (100 MW) * (1.0) * (18 Hours)
	+	(30 $/MWH) * (100 MW) * (1.0) * (70 Hours)
	+	(40 $/MWH) * (100 MW) * (1.0) * (70 Hours)
	+	(50 $/MWH) * (100 MW) * (1.0) * (10 Hours)
	=	$576,000 (34.29 $/MWH)
Unit 2:		(20 $/MWH) * (1000 MWH)
	+	(30 $/MWH) * (100 MW) * (1.0) * (70 Hours)

	+	(40 $/MWH) * (100 MW) * (1.0) * (70 Hours)
	+	(50 $/MWH) * (100 MW) * (1.0) * (10 Hours)
	=	$560,000 (35.00 $/MWH)
Unit 3:		(30 $/MWH) * (3800 MWH)
	+	(40 $/MWH) * (100 MW) * (1.0) * (70 Hours)
	+	(50 $/MWH) * (100 MW) * (1.0) * (10 Hours)
	=	$444,000 (37.63 $/MWH)
Unit 4:		(40 $/MWH) * (5600 MWH)
	+	(50 $/MWH) * (100 MW) * (1.0) * (10 Hours)
	=	$274,000 (41.52 $/MWH)

For each unit, the first term is for the hours that the unit itself is the marginal unit, except unit 1, which in this case is never on the margin. Succeeding terms add revenue for hours that later units are on the margin. The last term is for hours that there is insufficient capacity. For example, unit 3 generates 3800 MWH when it is itself the marginal unit and it earns revenue at its own price of 30 $/MWH. Unit 4 is on the margin for 70 hours, so unit 3 generates 100 MW * 1.0 (availability) * 70 hours = 7000 MWH when unit 4 is on the margin and the revenue is priced, based on unit 4, at 40 $/MWH. There are 10 hours when there is no marginal unit (loss of load); unit 3 generates 100 * 1.0 * 10 = 1000 MWH and earns revenue at the unserved energy price of 50 $/MWH for that energy. Thus, the total revenue for unit 3 is $444,000, which represents an average of 37.63 $/MWH.

For spinning reserve, the first block of a unit may be loaded out of economic order. In such a case, the revenue calculations are modified as follows:

The spinning reserve block itself earns revenues based on its own bid price until the second block of that unit is dispatched. (If EGEAS used the bid prices of units loaded between the first and second blocks of the spinning reserve unit, which are cheaper than the price of the spinning reserve unit, it would be penalizing the unit that provided the spinning reserve). After the second block is loaded, the spinning reserve block earns revenues based on the bid prices of the later units, in the same manner as any other unit.

Units loaded before the spinning reserve block earn revenue for the hours that the spinning reserve block is on margin. The price used for those revenues is the bid price of the last thermal unit loaded prior to the spinning reserve block. (If EGEAS used the bid price of the spinning reserve block, those earlier units would benefit by receiving revenue based on the high price of the unit moved earlier for spinning reserve).

For revenue calculations, limited energy plants (hydro units and the generating side of storage units) are treated in much the same manner as thermal units, with one important distinction: the bid price of the limited energy unit is not used. Instead, EGEAS uses the bid price of the thermal unit which is “split” by the limited energy unit. Thus, when one or more limited energy units are loaded with a thermal unit, EGEAS uses the bid price of the thermal unit for all the hours that the thermal unit and the limited energy units are marginal. The reason for this is that a limited energy unit can never be marginal – it is already generating the maximum amount of energy possible.

The charging sides of storage units add to the amount of generation that must be produced by thermal units. EGEAS assumes that a thermal generator that provides charging energy receives revenue at its own (i.e. the thermal units) bid price.

Non-dispatchable technologies (NDT's) are not dispatched according to economics as are the other types of units. Thus, an NDT unit is never the marginal unit and EGEAS never calculates revenues for NDT units. However, the user can input annual revenues for an NDT unit directly. Revenues can also be changed yearly with a trajectory.

Limitations

Storage units: if there are any storage units, either existing, committed, or planning alternatives, then storage projects must be defined in CANAL.

Storage projects: revenue calculations are made assuming deterministic allocation, even if the probabilistic allocation option is used for production costing.

Multi-area modeling: revenue calculations are not adjusted to reflect tie-line limits.

Fuel and emission constraints: when there are storage projects, such constraints are applied only in an approximate manner for the revenue calculations.

Renewable Portfolio Standard (RPS) Constraints

Renewable Portfolio Standards (RPS) are regulations designed to help increase the amount of renewable energy used for electric generation. The RPS constraints in EGEAS specify lower limits on the amount of energy which must be supplied to a generating system by certain renewable resources. These limits are specified in terms of percentages of the net system generation after energy efficiency programs have been factored in. Multiple RPS constraints, for different technologies or groupings of units, can be applied simultaneously. Constraints can be applied either with dynamic programming or prespecified pathway; constraints are not available with Benders decomposition. Limits can be changed from year to year.

RPS Calculation

The calculation of the RPS percentage considers energy provided to charge storage units, DSM energy, and energy from any other technologies that need special treatment, as well as the generation by RPS-eligible units. The user can specify that only a portion of a renewable unit's generation be counted in the RPS total. The RPS percentage is calculated as follows:

RPS % = [∑ RPS generation * rc] / [System energy + Charging – DSM – Excluded]

Where:
- RPS % = % generation from RPS (renewable) resources
- RPS generation = generation by an RPS unit, in Gwh
- rc = portion of an RPS unit's generation that is eligible to contribute to the RPS constraint
- System energy = total system energy requirements, in Gwh
- Charging = energy used to charge storage units, in Gwh

	DSM	= energy provided by demand-side management programs, in Gwh
	Excluded	= energy provided by other units designated as excluded (or netted out) from the calculation, in Gwh

Here is a sample calculation:

System energy	1,000 Gwh
Storage charging	50 Gwh
DSM energy	25 Gwh
Excluded units	75 Gwh
RPS (renewables)	150 Gwh

RPS % = 150 / [1,000 + 50 – 25 – 75] = 0.1579 or 15.79%

RPS Constraint Application

RPS constraints are applied at the same level of detail as production costing is performed - either annually or for each subperiod. After the production costing is completed, the program calculates the RPS percentage for each RPS constraint. If the RPS % falls below the minimum percentage for any constraint, that constraint is marked as not satisfied. For a dynamic program optimization, the current state is rejected as infeasible. For a prespecified pathway, an error message is generated, indicating that the constraint was not met.

The Electric Power Research Institute, Inc. (EPRI, www.epri.com) conducts research and development relating to the generation, delivery and use of electricity for the benefit of the public. An independent, nonprofit organization, EPRI brings together its scientists and engineers as well as experts from academia and industry to help address challenges in electricity, including reliability, efficiency, affordability, health, safety and the environment. EPRI also provides technology, policy and economic analyses to drive long-range research and development planning, and supports research in emerging technologies. EPRI's members represent approximately 90 percent of the electricity generated and delivered in the United States, and international participation extends to more than 30 countries. EPRI's principal offices and laboratories are located in Palo Alto, Calif.; Charlotte, N.C.; Knoxville, Tenn.; and Lenox, Mass.

Together…Shaping the Future of Electricity

© 2020 Electric Power Research Institute (EPRI), Inc. All rights reserved.
Electric Power Research Institute, EPRI, and TOGETHER…SHAPING THE FUTURE OF ELECTRICITY are registered service marks of the Electric Power Research Institute, Inc.

3002014878

Electric Power Research Institute
3420 Hillview Avenue, Palo Alto, California 94304-1338 • PO Box 10412, Palo Alto, California 94303-0813 • USA
800.313.3774 • 650.855.2121 • askepri@epri.com • www.epri.com